DEPARTMENT OF VETERANS AFFAIRS

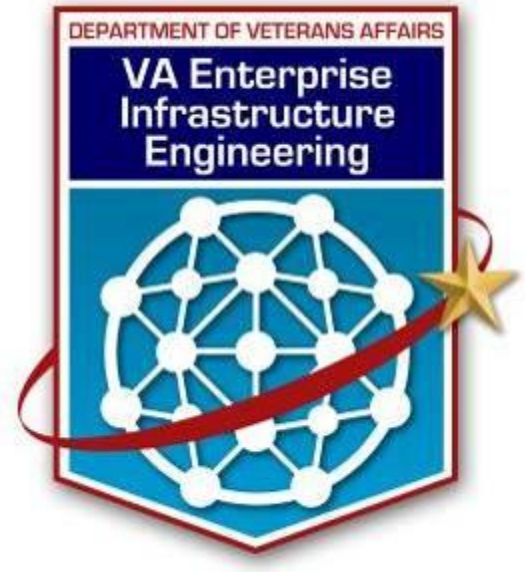

OFFICE OF INFORMATION AND TECHNOLOGY

ENTERPRISE INFRASTRUCTURE ENGINEERING

VA Enterprise IT Infrastructure Standard

Server Platform

PRODUCTION V1.0

November 18, 2009

TABLE OF CONTENTS

1 INTRODUCTION

1.1 PURPOSE

A standard is a set of rules or requirements that are determined by a consensus opinion of subject matter experts and prescribe criteria for a product, process, test or procedure. The general benefits of a standard are quality, interchangeability of parts or systems, and consistency. Information Technology (IT) standards are based on business needs provided through or supported by IT Services. IT Services are designed to support business processes and are constructed from software, hardware and infrastructure components. Establishing and enforcing standards for the selection and configuration of these supporting components improves the maintainability, reliability and availability of IT Services within projected economic constraints in alignment with business needs.

This standards document lists the acceptable and recommended specifications for Server Platforms. Sections include standard specifications for subject components, decisions supporting the standard specifications, guidelines or recommendations for implementing the standard specifications, and supplemental factors to consider for when evaluating subject components. Other supplementary documents will provide guidance on procuring components that meet the standard specifications, guidance on integrating them with existing components, and explanation of how the subject components fit into the VA Architecture.

1.2 OBJECTIVES

This standard provides acceptable and recommended specifications to support:

- Solution Evaluation
- Requirement Evaluation
- Solution Design
- Solution Procurement and Bid Evaluation
- Evaluation of Architectural Specifications
- Click here to enter Specific Objectives

1.3 SCOPE

This standard applies to:

- Specifications will include all attributes necessary to specify x86 class servers:
    - Server Hardware
    - Operating System
    - Supporting Equipment Rack Infrastructure

- All solutions that seek inclusion to the VA infrastructure undergo a technical analysis review (TAR). The TAR process is designed to ensure that all solutions meet the standards defined herein. A waiver can be granted if the solution can not meet the standards. Requests for waivers must go through the TAR process.
- To understand the TAR/TAS process, please refer to the following publications:
    - Anatomy of TAR Email Request.pdf
    - TAR-TAS Process-Steps.pdf
    - TAR Request Form - Project Name.docx

1.4 STANDARDS CLASSIFICATIONS

The server standards are divided into classes which encompass the needs of most current and future systems. The goal of the server classes is to simplify the platform selection process and to gain efficiencies in management of the platform. Having fewer standardized server platforms enables more efficient provisioning or re-provisioning of systems, the ability to have a test lab that accurately represents production options, and the ability to pre-buy systems enabling more rapid response to customer requests for new or increased capacity.

The server classes are derived through the analysis of server functions and workloads. Server functions are organized into four categories: content delivery, application server, transactional, and virtualization. To each of these categories is assigned three workloads: light, typical, and heavy. An evaluation of server function and workload exposes commonality among server requirements which leads to a concise set of server classes or sets of standards that can be used to accommodate most solution designs.

General Description of Server Functions:

Content Delivery – Examples: web server, FTP, file server, Active Directory / Domain Controller, SharePoint web portal, management server, print servers, DHCP servers

Application Server – Examples: WebLogic server, VistA front-end server, Terminal Server, SharePoint application tier server

Transactional – Examples: Database server, SharePoint MOSS server, data backup server

Virtualized Host – Examples: virtualization farm, virtualized content, application, or transactional servers

General Description of Server Classes:

Class A, the largest of the classes, targets transactional servers and virtualized host platforms with heavy workloads. Class B, having less performance and I/O capabilities than Class A, targets transactional servers with light to typical workloads and application and content delivery servers with heavy workloads. Class C is designed for content delivery servers with light to typical workloads. It should be noted that Class C servers still maintain significant performance characteristics and offer a smaller form factor then do class A and B servers. The class D specification has been set aside for a future standard which addresses cloud computing designs. Class E servers are transactional servers based on the Itanium processor and are used exclusively for VistA back-end (VistA BE) database servers. Class E servers are an exception to typical transactional servers defined by classes A and B.

Class A – VM hosts, heavy workload transactional

Class B –Heavy workload application servers, Heavy workload content delivery, light and typical workload transactional

Class C – Light and typical workload application servers, Light and typical workload content delivery, web servers

Class D – This class is based on cloud computing environments and has yet to be defined.

Class E – Typical to heavy workload Itanium based servers used exclusively for VistA BE designs.

The advantages and disadvantages of consolidating server types into fewer sets of defined categories, each with the highest "minimum" configuration required to serve the multiple roles, are:

Advantages:

- Allows easy structure of IDIQ contract
- Quicker to document
- Cost of buying extra memory, interfaces, is relatively small compared to overall cost and is less than trying to add those resources later through a server upgrade
- Simpler document to use for the audience/consumer when architecting or purchasing solutions
- Easier to re-provision for different workload or to modify workload on servers
- Easier to build/re-build and model test environments with a smaller variation of server types, and less chance of using "approximations" of server types in the test environments which can lead to higher quality of test results
- Creates a more converged standardization of servers in the data center for easier management and operations
- Easier to "pre-buy" the servers to set up a "utility" type of data center resource

Disadvantages:

- Less cost-efficient for initial (capital) acquisition costs
- Lower efficiency in operational utilization of resources, leading to more power consumption
- Potential for over-consumption of rack space (and subsequently floor space, more power circuits, more network infrastructure capacity assigned to racks)

How to Apply these Standards:

Server classes appear in this document in the following order: C, B, A, E. With the exception of the Class E server, the server classes are listed in order of workload capability from the lightest to the heaviest. The requirements of the Class E server represent an exception to the core standards set forth in this document and have been placed at the end.

Starting with the Class C server, each class adds to the previous class' characteristics in a way that increases its workload capability. The Class C server represents the least common denominator of all server classes – possessing the characteristics that fulfill many general purpose computing needs. These characteristics are refined as the workload requirements of the server increase, so that the Class B server has all of the Class C characteristics plus additional components (memory, processing capacity, etc.) to support its intended use. The Class A server, the largest server class, has all the Class B characteristics plus the additional components (again, largely memory, processing capacity, network connections, etc.) for its workload requirements.

As a general rule, servers should be procured to meet the requirements of their proposed workload. When applying these standards in a procurement process, the reader should first identify the function the server will perform and the workload the function will place on the server. Appendix F provides a list of common server functions and workloads and associates them with recommended server classes

It should be noted that for all classes, **the standards are the baseline specifications** and add-on requirements/enhancements can be used for system design and/or procurement – ex: additional NICs or RAM for certain workloads.

*** See Appendix F for table of server classes.**

2 STANDARDS

2.1 CLASS C SERVER

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
1	Processor	Type	x86_64 (e.g., Intel Xeon, AMD Opteron)
		Socket Count	Minimum required to meet Performance Benchmark standard
		Core Count	Amount required to meet Performance Benchmark standard
		Cache	Minimum total 2MB L2 Cache on chip
		Virtualization Aware	Not Required
2	Performance	Performance Benchmark	SPEC CPU2006 Integer Rate Base value >= 175
3	Memory	Type	Compatible with CPU with full error correction capability including ECC; all memory modules must be of the same type and size.
		Total Memory	Minimum 16GB usable
		Speed	Minimum – match operational CPU and bus speed required to meet the Performance Benchmark requirement
4	Internal Storage	Type	SAS RAID controller supporting SAS or SATA drives and RAID/hot-spare features specified for this server
		Quantity of Drives	Minimum three
		RAID Configuration	Minimally require 2 disks in RAID1, plus 1 as hot spare, for page/swap, dump file, and options operating system functional role. Additional local storage needs (e.g., "temp file" storage for a transaction server or standalone storage configurations) are satisfied by additional drives in RAID level 1, 5, or 10 configuration, with a hot spare which can be shared with the page/swap disk's hot spare.
		Drive Speed	Minimum 10K RPM
		Drive Capacity	Minimum 70GB usable storage capacity
		Disk Subsystem Performance	The disk raid group must meet or exceed the performance levels specified for tests "Transaction Rate" and "Throughput Rate" Validation will be incorporated in

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
			user acceptance phase.
		Controller Cache	Minimum 256MB
		Quantity of Controllers	Minimum of one
5	External Storage Interface	Type	Fiber-channel with duplex LC connectors, CPU bus type is PCI-E Or Ethernet matching connector type of storage-dedicated network in data center
		Port Count	Minimum of two
		Bandwidth	Minimum 4Gbps for optical or 1Gbps for copper; must meet or exceed SAN fabric requirements
		Compatibility	Support storage arrays at deployment location; supported by operating system and multipath drivers specified for this server
		SAN Boot Capability	BIOS support for SAN boot
		Blade Chassis Component (if applicable)	Dual embedded switches
		Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the SAN; minimum redundant ports. Support for link aggregation.
6	Server Communications Interface	Type	Ethernet matching connector type for network in data center
		Port Count	Minimum of 2 ports; at least 2 more if private cluster interconnect is required
		Bandwidth	Minimum 1Gbps
		Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the LAN; supports redundant ports (may include embedded ports) for each interface function (e.g., public network function and cluster interconnect function)
		Blade Chassis Component (if applicable)	Dual embedded switches

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
7	Removable Media Device	Optical Drive	(optional) CD/DVD drive – minimum 24x speed
		USB	USB v2.0 port
8	Power	Redundancy	N+1 redundancy on power supply
		Voltage	115-208V input voltage on power supplies
		Cord Type	IEC
		Power Efficiency	Until such time that there are Energy Star ratings available to most server classes, power efficiency should be consistent with industry trends.
		Power Consumption Documentation	Vendor-supplied data for power consumption of specified configuration at 100% load, “typical” active usage, and at idle.
9	Cooling Fan	Redundancy	Minimum N+1 redundancy
10	Out of Band Management	Type	Remote command-line and console-level access utilizing dedicated network interface
		Security	Secure IP-based remote management that complies with all VA security requirements
		Remote Power Control	Ability to remotely power on/off/reset server
11	Server Hardware Vendor Support	Type	Match service level requirement for hosted applications; mission-critical level support shall include 24x7 coverage 365 days/year, less than 1 hour initial engineer-level response time, 4-hour on-site response time for emergency dispatches, and 8-hour time-to-repair
		Term Length	Minimum of three years with two additional option years
		Keep Your Hard Drive	All storage drives are kept by VA in the event of warranty/service replacement
12	Hardware Management Software	Type	Agent or agent-less server management tool accessible through SNMP
		Automated Notification	Hardware events including exceptions, diagnostics, and failures, shall be exposed through SNMP with documented MIBs
		Shared Repository	Allows uploading all information collected by the tool into shared repository using ODBC or JBDC in a manner

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
			that preserves data fields for the diagnostic information (e.g., date/time, error code, module, description, system identification, status, system configuration) Refer to Datacenter Standards, section _____(name, no numbers)
		Support Level	Support level of management software must match support level of server hardware
13	Rack Infrastructure Compatibility	Rack Unit Measurement	1U or 2U for rack-mount servers; blade implementations will be evaluated on an individual basis.
		Rail Type	Tool-less square-hole sliding
		Cable Management	Side-reversible for non-blade solutions
14	Operating System	Type	Linux preferred based on federal requirement to utilize open source operating systems, Windows acceptable, all must VA approved operating systems and compatible with hardware.
		Version	Current target version of Linux or Windows per VA TRM
		Vendor-Installed	Depends on data center policy; default is factory-installed with default settings and most recent patch level
		Support Level / Term	Must match hardware service level; term is three years with two option years unless VA has a standing enterprise software support agreement pertinent OS.
15	Platform Software / License	Multipath	Provides automated load balancing (all paths active) and fail-over; compatible with target storage array
		Data Backup	(optional) Matches requirements of the deployment location data center and SAN backup solution
		Monitoring Tool	Matches requirements of the deployment location data center
		Layered Framework	Defined by the application owner – meet license requirements for this server
		Server Management Tools	Matches requirements of the deployment location data center
		Availability / Recovery Tools	Matches the requirements for Persistent Computing standards applied to the hosted applications/services
		Security Tools	Follows Office of Information and Technology standards

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
			including compliance with FIPS, FISMA, and VA Directives.
		Virtual Management Software	Not Required
16	Datacenter Management	Remote Keyboard/Video/Mouse Control	Ability to remotely connect to the server's external keyboard, video, and mouse ports through a networked KVM switch.
		Remote Power Disconnect	Ability to remotely disconnect/reconnect server to power source

2.2 CLASS B SERVER

[Populate this table from the taxonomy in section 4 and enter specifications as necessary. Replace this paragraph with a brief description of this specification set.

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
1	Processor	Type	x86_64 (e.g., Intel Xeon, AMD Opteron)
		Socket Count	Minimum required to meet Performance Benchmark standard
		Core Count	Amount required to meet Performance Benchmark standard
		Cache	Minimum total 2MB L2 Cache on chip (6MB or greater L3 Cache required for transactional servers)
		Virtualization Aware	Not Required
2	Performance	Performance Benchmark	SPEC CPU2006 Integer Rate Base value >= 200
3	Memory	Type	Compatible with CPU with full error correction capability including ECC; all memory modules must be of the same type and size.
		Total Memory	Minimum 48GB usable
		Speed	Minimum – match operational CPU and bus speed required to meet the Performance Benchmark requirement
4	Internal Storage	Type	SAS RAID controller supporting SAS drives and RAID/hot-spare features specified for this server
		Quantity of Drives	Minimum three
		RAID Configuration	Minimally require 2 disks in RAID1, plus 1 as hot spare, for page/swap, dump file, and options operating system functional role. Additional local storage needs (e.g., “temp file” storage for a transaction server or standalone storage configurations) are satisfied by additional drives in RAID level 1, 5, or 10 configuration, with a hot spare which can be shared with the page/swap disk’s hot spare.
		Drive Speed	Minimum 10K RPM
		Drive Capacity	Minimum 300GB usable storage capacity

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
		Disk Subsystem Performance	The disk raid group must meet or exceed the performance levels specified for tests “Transaction Rate” and “Throughput Rate” Validation will be incorporated in user acceptance phase.
		Controller Cache	Minimum 256MB
		Quantity of Controllers	Minimum of one
5	External Storage Interface	Type	Fiber-channel with duplex LC connectors, CPU bus type is PCI-E Or Ethernet matching connector type of storage-dedicated network in data center
		Port Count	Minimum of two
		Bandwidth	Minimum 4Gbps for optical or 1Gbps for copper; must meet or exceed SAN fabric requirements
		Compatibility	Support storage arrays at deployment location; supported by operating system and multipath drivers specified for this server
		SAN Boot Capability	BIOS support for SAN boot
		Blade Chassis Component (if applicable)	Dual embedded switches
		Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the SAN. Minimum two adapters with minimum one port per adapter. Support for link aggregation.
6	Server Communications Interface	Type	Ethernet matching connector type for network in data center
		Port Count	Minimum of 2 ports; at least 2 more if private cluster interconnect is required
		Bandwidth	Minimum 1Gbps
		Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the LAN; supports redundant ports (may include embedded ports) for each interface function (e.g., public network function and cluster

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
			interconnect function)
		Blade Chassis Component (if applicable)	Dual embedded switches
7	Removable Media Device	Optical Drive	(optional) CD/DVD drive – minimum 24x speed
		USB	USB v2.0 port
8	Power	Redundancy	N+1 redundancy on power supply
		Voltage	115-208V input voltage on power supplies
		Cord Type	IEC
		Power Efficiency	Until such time that there are Energy Star ratings available to most server classes, power efficiency should be consistent with industry trends.
		Power Consumption Documentation	Vendor-supplied data for power consumption of specified configuration at 100% load, “typical” active usage, and at idle.
9	Cooling Fan	Redundancy	Minimum N+1 redundancy
10	Out of Band Management	Type	Remote command-line and console-level access utilizing dedicated network interface
		Security	Secure IP-based remote management that complies with all VA security requirements
		Remote Power Control	Ability to remotely power on/off/reset server
11	Server Hardware Vendor Support	Type	Match service level requirement for hosted applications; mission-critical level support shall include 24x7 coverage 365 days/year, less than 1 hour initial engineer-level response time, 4-hour on-site response time for emergency dispatches, and 8-hour time-to-repair
		Term Length	Minimum of three years with two additional option years
		Keep Your Hard Drive	All storage drives are kept by VA in the event of warranty/service replacement
12	Hardware Management Software	Type	Agent or agent-less server management tool accessible through SNMP
		Automated Notification	Hardware events including exceptions, diagnostics, and failures, shall be exposed through SNMP with

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
			documented MIBs
		Shared Repository	Allows uploading all information collected by the tool into shared repository using ODBC or JBDC in a manner that preserves data fields for the diagnostic information (e.g., date/time, error code, module, description, system identification, status, system configuration)
		Support Level	Support level of management software must match support level of server hardware
13	Rack Infrastructure Compatibility	Rack Unit Measurement	1U or 2U for rack-mount servers; blade implementations will be evaluated on an individual basis.
		Rail Type	Tool-less square-hole sliding
		Cable Management	Side-reversible for non-blade solutions
14	Operating System	Type	Linux preferred based on federal requirement to utilize open source operating systems, Windows acceptable, all must VA approved operating systems and compatible with hardware.
		Version	Current target version of Linux or Windows per VA TRM
		Vendor-Installed	Depends on data center policy; default is factory-installed with default settings and most recent patch level
		Support Level / Term	Must match hardware service level; term is three years with two option years unless VA has a standing enterprise software support agreement pertinent OS.
15	Platform Software / License	Multipath	Provides automated load balancing (all paths active) and fail-over; compatible with target storage array
		Data Backup	(optional) Matches requirements of the deployment location data center and SAN backup solution
		Monitoring Tool	Matches requirements of the deployment location data center
		Layered Framework	Defined by the application owner – meet license requirements for this server
		Server Management Tools	Matches requirements of the deployment location data center
		Availability / Recovery Tools	Matches the requirements for Persistent Computing

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
			standards applied to the hosted applications/services
		Security Tools	Follows Office of Information and Technology standards including compliance with FIPS, FISMA, and VA Directives.
		Virtual Management Software	Not Required
16	Datacenter Management	Remote Keyboard/Video/Mouse Control	Ability to remotely connect to the server's external keyboard, video, and mouse ports through a networked KVM switch.
		Remote Power Disconnect	Ability to remotely disconnect/reconnect server to power source

2.3 CLASS A SERVER

[Populate this table from the taxonomy in section 4 and enter specifications as necessary. Replace this paragraph with a brief description of this specification set.

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
1	Processor	Type	x86_64 (e.g., Intel Xeon, AMD Opteron)
		Socket Count	Minimum required to meet Performance Benchmark standard
		Core Count	Amount required to meet Performance Benchmark standard
		Cache	Minimum total 2MB L2 Cache on chip (6MB or greater L3 Cache required for transactional servers)
		Virtualization Aware	Required to facilitate live migration
2	Performance	Performance Benchmark	SPEC CPU2006 Integer Rate Base value >= 250
3	Memory	Type	Compatible with CPU with full error correction capability including ECC; all memory modules must be of the same type and size.
		Total Memory	Minimum 128GB usable
		Speed	Minimum - match operational CPU and bus speed required to meet the Performance Benchmark requirement
4	Internal Storage	Type	SAS RAID controller supporting SAS drives and RAID/hot-spare features specified for this server
		Quantity of Drives	Minimum three
		RAID Configuration	Minimally require 2 disks in RAID1, plus 1 as hot spare, for page/swap, dump file, and options operating system functional role. Additional local storage needs (e.g., "temp file" storage for a transaction server or standalone storage configurations) are satisfied by additional drives in RAID level 1, 5, or 10 configuration, with a hot spare which can be shared with the page/swap disk's hot spare. When this class of server is configured as a transaction

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
			server, a separate RAID controller or second channel on the same controller is required for the “temp file” RAID group
		Drive Speed	Minimum 10K RPM
		Drive Capacity	Minimum 450GB usable storage capacity
		Disk Subsystem Performance	The disk raid group must meet or exceed the performance levels specified for tests “Transaction Rate” and “Throughput Rate” Validation will be incorporated in user acceptance phase.
		Controller Cache	Minimum 256MB
		Quantity of Controllers	Minimum of one
5	External Storage Interface	Type	Fiber-channel with duplex LC connectors, CPU bus type is PCI-E Or Ethernet matching connector type of storage-dedicated network in data center
		Port Count	Minimum of 2 independent cards with a minimum of 2 ports each
		Bandwidth	Minimum 4Gbps for optical or 1Gbps for copper; must meet or exceed SAN fabric requirements
		Compatibility	Support storage arrays at deployment location; supported by operating system and multipath drivers specified for this server
		SAN Boot Capability	BIOS support for SAN boot
		Blade Chassis Component (if applicable)	Dual embedded switches
		Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the SAN. Minimum two adapters with minimum one port per adapter. Support for link aggregation.
6	Server Communications Interface	Type	Ethernet matching connector type for network in data center
		Port Count	Minimum of 6 ports; more may be necessary depending

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
			on virtual infrastructure requirements
		Bandwidth	Minimum 1Gbps
		Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the LAN; supports physically independent redundant ports including embedded ports for each interface function (e.g., public network function and cluster interconnect function)
		Blade Chassis Component (if applicable)	Dual embedded switches
7	Removable Media Device	Optical Drive	(optional) CD/DVD drive – minimum 24x speed
		USB	USB v2.0 port
8	Power	Redundancy	N+1 redundancy on power supply
		Voltage	208V input voltage on power supply
		Cord Type	IEC
		Power Efficiency	Until such time that there are Energy Star ratings available to most server classes, power efficiency should be consistent with industry trends.
		Power Consumption Documentation	Vendor-supplied data for power consumption of specified configuration at 100% load, “typical” active usage, and at idle.
9	Cooling Fan	Redundancy	Minimum N+1 redundancy
10	Out of Band Management	Type	Remote command-line and console-level access utilizing dedicated network interface
		Security	Secure IP-based remote management that complies with all VA security requirements
		Remote Power Control	Ability to remotely power on/off/reset server
11	Server Hardware Vendor Support	Type	Match service level requirement for hosted applications; mission-critical level support shall include 24x7 coverage 365 days/year, less than 1 hour initial engineer-level response time, 4-hour on-site response time for emergency dispatches, and 8-hour time-to-repair
		Term Length	Minimum of three years with two additional option years

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
		Keep Your Hard Drive	All storage drives are kept by VA in the event of warranty/service replacement
12	Hardware Management Software	Type	Agent or agent-less server management tool accessible through SNMP
		Automated Notification	Hardware events including exceptions, diagnostics, and failures, shall be exposed through SNMP with documented MIBs
		Shared Repository	Allows uploading all information collected by the tool into shared repository using ODBC or JBDC in a manner that preserves data fields for the diagnostic information (e.g., date/time, error code, module, description, system identification, status, system configuration)
		Support Level	Support level of management software must match support level of server hardware
13	Rack Infrastructure Compatibility	Rack Unit Measurement	Up to 4U for rack-mount servers; blade implementations will be evaluated on an individual basis.
		Rail Type	Tool-less square-hole sliding
		Cable Management	Side-reversible for non-blade solutions
14	Operating System	Type	Linux preferred based on federal requirement to utilize open source operating systems, Windows acceptable, all must VA approved operating systems and compatible with hardware.
		Version	Current target version of Linux or Windows per VA TRM
		Vendor-Installed	Depends on data center policy; default is factory-installed with default settings and most recent patch level
		Support Level / Term	Must match hardware service level; term is three years with two option years unless VA has a standing enterprise software support agreement pertinent OS.
15	Platform Software / License	Multipath	Provides automated load balancing (all paths active) and fail-over; compatible with target storage array
		Data Backup	(optional) Matches requirements of the deployment location data center and SAN backup solution
		Monitoring Tool	Matches requirements of the deployment location data center

ID	Primary Attribute	Secondary Attribute	Specification
		Layered Framework	Defined by the application owner – meet license requirements for this server
		Server Management Tools	Matches requirements of the deployment location data center
		Availability / Recovery Tools	Matches the requirements for Persistent Computing standards applied to the hosted applications/services
		Security Tools	Follows Office of Information and Technology standards including compliance with FIPS, FISMA, and VA Directives.
		Virtual Management Software	For virtual environments, an Enterprise management tool is required
16	Datacenter Management	Remote Keyboard/Video/Mouse Control	Ability to remotely connect to the server's external keyboard, video, and mouse ports through a networked KVM switch.
		Remote Power Disconnect	Ability to remotely disconnect/reconnect server to power source

2.4 CLASS E SERVER

[Populate this table from the taxonomy in section 4 and enter specifications as necessary. Replace this paragraph with a brief description of this specification set.

ID	Primary Attribute	Secondary Attribute	Specification
1	Processor	Type	IA_64 (Itanium)
		Socket Count	Minimum required to meet Performance Benchmark standard
		Core Count	Amount required to meet Performance Benchmark standard
		Cache	Minimum required to meet Performance Benchmark standard
		Virtualization Aware	Not Required

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
2	Performance	Performance Benchmark	SPEC CPU2006 Integer Rate Base value >= 90
3	Memory	Type	Compatible with CPU with full error correction capability including ECC
		Total Memory	Minimum 24GB usable
		Speed	Minimum - match operational CPU and bus speed required to meet the Performance Benchmark requirement
4	Internal Storage	Type	SAS
		Quantity of Drives	Minimum one (DOSD – i.e. to host Dump Off System Disk)
		RAID Configuration	Not Required
		Drive Speed	Minimum 10K RPM
		Drive Capacity	Minimum 36GB usable storage capacity
		Disk Subsystem Performance	The disk must meet or exceed the performance levels specified below for tests “Transaction Rate” and “Throughput Rate” Validation will be incorporated in user acceptance phase.
		Controller Cache	Not Required
		Quantity of Controllers	Minimum of one
5	External Storage Interface	Type	Fiber-channel with duplex LC connectors
		Port Count	Minimum of four ports with full port redundancy provided.
		Bandwidth	Minimum 4Gbps; must meet or exceed SAN fabric requirements
		Compatibility	Support storage arrays at deployment location; supported by operating system and multipath drivers specified for this server
		SAN Boot Capability	BIOS support for SAN boot
		Blade Chassis Component (if applicable)	Pass thru connections, for increased bandwidth capabilities

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
		Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the SAN. Minimum redundant ports
6	Server Communications Interface	Type	Ethernet matching connector type for network in data center
		Port Count	Minimum of 6 ports
		Bandwidth	Minimum 1Gbps
		Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the LAN; supports physically independent redundant ports including embedded ports for each interface function (e.g., public network function and cluster interconnect function)
		Blade Chassis Component (if applicable)	Pass Thru connections
7	Removable Media Device	Optical Drive	CD/DVD drive – minimum 24x speed
		USB	Not Required
8	Power	Redundancy	N+1 redundancy on power supply
		Voltage	208V input voltage on power supply
		Cord Type	IEC
		Power Efficiency	Until such time that there are Energy Star ratings available to most server classes, power efficiency should be consistent with industry trends.
		Power Consumption Documentation	Vendor-supplied data for power consumption of specified configuration at 100% load, “typical” active usage, and at idle.
9	Cooling Fan	Redundancy	Minimum N+1 redundancy
10	Out of Band Management	Type	Remote command-line and console-level access utilizing dedicated network interface
		Security	Secure IP-based remote management that complies with all VA security requirements
		Remote Power Control	Ability to remotely power on/off/reset server

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
11	Server Hardware Vendor Support	Type	Match service level requirement for hosted applications; mission-critical level support shall include 24x7 coverage 365 days/year, less than 1 hour initial engineer-level response time, 4-hour on-site response time for emergency dispatches, and 8-hour time-to-repair
		Term Length	Minimum of three years with two additional option years
		Keep Your Hard Drive	All storage drives are kept by VA in the event of warranty/service replacement
12	Hardware Management Software	Type	Agent or agent-less server management tool accessible through SNMP
		Automated Notification	Hardware events including exceptions, diagnostics, and failures, shall be exposed through SNMP with documented MIBs
		Shared Repository	Allows uploading all information collected by the tool into shared repository using ODBC or JBDC in a manner that preserves data fields for the diagnostic information (e.g., date/time, error code, module, description, system identification, status, system configuration)
		Support Level	Support level of management software must match support level of server hardware
13	Rack Infrastructure Compatibility	Rack Unit Measurement	7U or 10U for rack-mount servers; 4-slot blade for blade servers
		Rail Type	Tool-less square-hole sliding
		Cable Management	Side-reversible for non-blade solutions
14	Operating System	Type	OpenVMS, Un-Limited User Licensing, VMS clustering, Host Based Volume Shadowing, TCP/IP Services
		Version	OpenVMS version 8.3 (minimum version support)
		Vendor-Installed	Not Required, but preferred
		Support Level / Term	Must match hardware service level; term is three years with two option years unless VA has a standing enterprise software support agreement pertinent OS.
15	Platform Software / License	Multipath	Provides automated load balancing (all paths active) and fail-over; compatible with target storage array

ID	*Primary Attribute*	*Secondary Attribute*	*Specification*
		Data Backup	(optional) Matches requirements of the deployment location data center and SAN backup solution
		Monitoring Tool	Matches requirements of the deployment location data center
		Layered Framework	Support for InterSystems' Caché (minimum version 5.2.3)
		Server Management Tools	Matches requirements of the deployment location data center
		Availability / Recovery Tools	Matches the requirements for Persistent Computing standards applied to the hosted applications/services
		Security Tools	Follows Office of Information and Technology standards including compliance with FIPS, FISMA, and VA Directives.
		Virtual Management Software	Not Required
16	Datacenter Management	Remote Keyboard/Video/Mouse Control	Ability to remotely connect to the server's external keyboard, video, and mouse ports through a networked KVM switch.
		Remote Power Disconnect	Ability to remotely disconnect/reconnect server to power source

3 SUPPORTING DETAILS FOR STANDARDS

3.1 LIGHT WORKLOAD **(CLASS C)**

3.1.1 PROCESSOR

STANDARD

ID	*Secondary Attribute*	*Specification*
1	Type	x86_64 (e.g., Intel Xeon, AMD Opteron)
	Socket Count	Minimum required to meet Performance Benchmark standard
	Core Count	Amount required to meet Performance Benchmark standard
	Cache	Minimum total 2MB L2 Cache on chip
	Virtualization Aware	Not Required

EXPLANATION OF STANDARD

Type – open/industry standard components.

Socket Count – minimize due to licensing costs per socket, have minimum count for performance/capacity.

Core Count – minimum amount expected for performance/capacity level is quad-core, more can be better. However, any amount that meets the Performance Benchmark specification is sufficient.

Cache – minimum expected for class of processor; low-power options might use lower amounts of cache.

EVALUATION FACTORS

- Currency of processor – most recent family and speed of processor is beneficial.
- Higher amounts of cache and cores are beneficial.
- Open processor slots or ability to enable additional processors for future expansion could be beneficial in larger configurations.

IMPLEMENTATION GUIDANCE

It is recommended to have a Socket Count of two in order to minimize licensing costs that are calculated on a per socket basis. Exceptions for higher socket count should have justification documented and should consider both performance and cost.

3.1.2 PERFORMANCE

STANDARD

ID	*Secondary Attribute*	*Specification*
2	Performance Benchmark	SPEC CPU2006 Integer Rate Base value >= 175

EXPLANATION OF STANDARD

This SPEC_Int rate is represented by a large number of server platforms. Performance requirement is largely dependent on the application's workload and compute-resource consumption profile. The architecture of this server tier may provide sufficient horizontal scaling to accommodate larger workloads, especially with minimum server counts required for redundancy and limit of single-server impact. This specification is a starting point for a "light duty" small-workload content delivery server environment based upon review of existing tested systems in the VA environment and their capabilities represented with newer server offerings. This specification should be adjusted to meet the known workload of target hosted applications and the target architecture for its server tier.

EVALUATION FACTORS

Higher benchmark rating value is beneficial.

IMPLEMENTATION GUIDANCE

3.1.3 MEMORY

STANDARD

ID	*Secondary Attribute*	*Specification*
3	Type	Compatible with CPU with full error correction capability including ECC; all memory modules must be of the same type and size.
	Total Memory	Minimum 16GB usable
	Speed	Minimum – match operational CPU and bus speed required to meet the Performance Benchmark requirement

EXPLANATION OF STANDARD

Type – Memory has to be highly reliable – ECC is the expected feature to allow higher levels of reliability. RDIMM is more expandable/scalable and provides additional reliability features (e.g., more error-checking features) while

consuming more power (about one watt per chip) and costing more. In general, RDIMM is recommended with high memory requirements and higher availability requirements.

Total Memory – Memory requirements depend on application needs. The standard must meet the minimum requirement for memory intensive applications that could be run on this class of server. Using WebLogic as an example, minimum expected is 2GB-4GB per CPU core (as per Java configuration guidelines). With current VA application designs, each JVM requires at least 4GB of memory; more if 64 bit JVMs are used. For a typical dual-socket quad-core server running WebLogic, the minimum expected requirement is therefore 16GB. Usable memory value indicates the possibility of RAID or mirrored memory for additional redundancy. When using memory configurations above the minimum, internal storage capacity must be adjusted to support page/swap and dump requirements. This should be calculated using the formula: (2XRAM=page/swap size) + (1XRAM=DUMP) + (20GB for OS)

Speed – At least one operational speed of the memory must match the highest speed capability of the information bus used to retrieve memory for the CPU in order to maximize the investment into the CPU. Features can allow for step-down of speed correlated with workload consumption (power-saving feature).

EVALUATION FACTORS

- Additional memory is beneficial
- Higher density chips while avoiding a decrease in memory bus speed is preferred.
- Speed of memory takes precedence over density as long as capacity is met

IMPLEMENTATION GUIDANCE

RDIMM is recommended for high memory requirements and higher availability requirements.

3.1.4 INTERNAL STORAGE

STANDARD

ID	*Secondary Attribute*	*Specification*
4	Type	SAS RAID controller supporting SAS or SATA drives and RAID/hot-spare features specified for this server
	Quantity of Drives	Minimum three
	RAID Configuration	Minimally require 2 disks in RAID1, plus 1 as hot spare, for page/swap, dump file, and options operating system functional role. Additional local storage needs (e.g., “temp file” storage for a transaction server or standalone storage configurations) are satisfied by additional drives in RAID level 1, 5, or 10 configuration, with a hot spare which can be shared with the page/swap disk’s hot spare.
	Drive Speed	Minimum 10K RPM
	Drive Capacity	Minimum 70GB usable storage capacity

	Disk Subsystem Performance	The disk raid group must meet or exceed the performance levels specified for tests “Transaction Rate” and “Throughput Rate” Validation will be incorporated in user acceptance phase.
	Controller Cache	Minimum 256MB
	Quantity of Controllers	Minimum of one

EXPLANATION OF STANDARD

Type – SAS or SATA drives may be used. While SAS is preferred, SATA may be used where capacity or cost takes precedence over performance. Solid state drives are beneficial due to power, cooling, performance, and reliability, but are expensive and currently are only used sparingly in external storage systems.

Quantity & RAID Configuration – requires performance for a page/swap disk and dump file, with possible additional uses; ability to continue with a drive failure (e.g., RAID1) and restore to this state with an immediately available drive (hot spare). For high performance functions (including page/swap disk), RAID1 is preferred over parity protection due to overhead of calculating parity. Although additional drives may help spread I/O across more drives for enhanced performance, additional drives also add to the cost, maintenance, and power consumption on the server. Therefore, additional drives should only be considered when they are required for performance or capacity reasons. The optimal configuration uses the fewest quantity of drives to meet the performance and reliability requirements. Additional drives are needed for additional storage needs such as a standalone storage configuration which is provided on a separate disk per industry best practice guidance.

Drive Speed – Minimum RPM for this class of server & storage is 10K.

Drive Capacity – 70GB minimum usable storage capacity is sufficient to support an operating system and/or swap and log files. This amount represents the expected minimum required for page/swap file, dump file, and optional system disk storage needs. When using memory configurations above the minimum, internal storage capacity must be adjusted to support page/swap and dump requirements. This should be calculated using the formula: (2XRAM=page/swap size) + (1XRAM=DUMP) + (20GB for OS)

Disk Subsystem Performance – The disk raid group must meet or exceed the performance levels specified below for tests “Transaction Rate” and “Throughput Rate” as described in APPENDIX G.

- Transaction Rate = 500+ IOPS
- Throughput Rate = 50+ MBPS

Controller Cache – minimum expected value for typical controllers for this class of server.

Quantity of Controllers – limited value to have more than one with only one RAID configuration; might be beneficial for more than two RAID sets if performance requires it; current technology offerings limit the ability to fail services over from one controller to another without manual intervention – future technology offerings such as dual-domain active/active configurations in SAS controllers & drives provide this capability

EVALUATION FACTORS

- Higher storage capacity is beneficial
- Higher RPM is beneficial for drive platters
- The optimal configuration uses the fewest quantity of drives to meet the performance and reliability requirements.
- Solid-state drives would be beneficial
- Higher throughput capability is beneficial
- Higher IOPS capability is beneficial
- More controllers can be beneficial
- More cache is beneficial
- Battery backup of cache is beneficial
- Support for RAID levels 0, 1, 5, 6, 0+1, and/or 1+0 is beneficial
- Separate hot spares for separate RAID groups is beneficial.

IMPLEMENTATION GUIDANCE

- SAS interface type is recommended. Battery backup of cache is recommended.

3.1.5 EXTERNAL STORAGE INTERFACE

STANDARD

ID	*Secondary Attribute*	*Specification*
5	Type	Fiber-channel with duplex LC connectors, CPU bus type is PCI-E Or Ethernet matching connector type of storage-dedicated network in data center. TOE / TCP checksum offload required.
	Port Count	Minimum of two
	Bandwidth	Minimum 4Gbps for optical or 1Gbps for copper; must meet or exceed SAN fabric requirements
	Compatibility	Support storage arrays at deployment location; supported by operating system and multipath drivers specified for this server
	SAN Boot Capability	BIOS support for SAN boot
	Blade Chassis Component (if applicable)	Dual embedded switches

	Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the SAN; minimum redundant ports. Support for link aggregation.

EXPLANATION OF STANDARD

- Type – fiber channel and Ethernet are the vast majority of investment for connections used in storage-dedicated networks in VA data centers, and exceptions should be justified including impact to existing infrastructure investment. Typical Ethernet connector types are RJ-45 for copper and LC for optical. When Ethernet is used, TOE / TCP checksum offload is required. Drivers are more readily available for Windows than Linux servers.
- Port Count – minimum required for port-level redundancy / load-balancing
- Bandwidth – 4Gbps is standard for current investment of optical fabrics in the storage systems and infrastructure in VA data centers and matches the minimum current industry standard; similarly 1Gbps is standard for current investments and minimum current industry practices when using copper-based storage fabrics.
- Compatibility – Support storage subsystems at deployment locations. Check adapter compatibility for storage at deployment location.
- Blade Chassis Component – A primary benefit of blade servers is their ability to reduce datacenter cabling. Dual fiber-channel switches (or Ethernet if applicable) are required. Pass-thru configurations require an exception. Each blade chassis shall have redundant sets of embedded switches to provide independent SAN connections to each blade. Each embedded switch set shall have the ability to configure redundant and/or trunked ISL/uplink ports. Minimum number of uplink ports in chassis is two times the number of embedded switches. Refer to Storage Standards, Section xx.
- Hardware Redundancy – Fail-over and load-balancing are functions of the multipath tool which is covered in section 3.2.16

EVALUATION FACTORS

- Redundancy across physically independent interfaces is beneficial
- Support for 8Gbps optical or 10Gbps copper bandwidth is beneficial

IMPLEMENTATION GUIDANCE

- For a server with high criticality of SAN storage connectivity, redundant adapter cards are recommended.
- Interface for iSCSI can be termed an HBA interface or a NIC interface. Where higher performance is a concern, the HBA interface is recommended.
- Bandwidth spread across multiple ports on independent cards may be required for performance reasons

3.1.6 SERVER COMMUNICATIONS INTERFACE

STANDARD

ID	*Secondary Attribute*	*Specification*

6	Type	Ethernet matching connector type for network in data center
	Port Count	Minimum of 2 ports; at least 2 more if private cluster interconnect is required
	Bandwidth	Minimum 1Gbps
	Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the LAN; supports redundant ports (may include embedded ports) for each interface function (e.g., public network function and cluster interconnect function)
	Blade Chassis Component (if applicable)	Dual embedded switches

EXPLANATION OF STANDARD

Type – Ethernet is the vast majority of investment for connections used in data center networks in VA, and exceptions should be justified including impact to existing infrastructure investment. Typical Ethernet connector types are RJ-45 for copper and LC for optical with a future trend towards copper only.

Port Count – base set of 2 ports are for general "public" Ethernet; other specific interfaces may be required depending on the profile of the workload/application. Cluster interconnects may need 2 or more ports, depending on the amount of cluster traffic (e.g., Oracle clusters send data over the cluster interconnect, simple Microsoft clustering only send "heartbeat" traffic).

Bandwidth – expected standard interface on network switches is 1Gbps. There is a potential for 10Gbps Ethernet.

Hardware Redundancy – embedded ports count as one "card" which is sufficient for this class of server except where additional ports are required. If additional ports are required, all interface uses are spread across two "cards" (including at least one add-on NIC). The goal is physical separation of the chipset/ASIC used on the redundant port, for each function. Support for "simple failover" (with a single IP address) and independent usage (with multiple IP addresses) is required.

Blade Chassis Component – Where applicable, blade-server chassis should be configured with two switches to limit port and cabling needs external to the server, unless there is a specific workload/application requirement for high throughput that requires pass-through capability.

EVALUATION FACTORS

- Additional cards can be beneficial
- Support for TCP Offload Engine is beneficial
- Support for active/active load balancing is beneficial

IMPLEMENTATION GUIDANCE

- One or two additional ports are recommended for network backups if the impact of backup traffic may affect other services on the "public" network. One port is sufficient if only one port is active and uses another port for fail-over.

- TCP Offload Engine (TOE) is preferred for all servers. Drivers are more readily available for Windows than Linux servers.
- If there is a specific workload/application requirement for high throughput that requires pass-through capability on a blade-server chassis, then this option should be configured, but the chassis switch is preferred otherwise.

BANDWIDTH SPREAD ACROSS MULTIPLE PORTS ON INDEPENDENT CARDS MAY BE REQUIRED FOR PERFORMANCE REASONS

3.1.7 REMOVABLE MEDIA DEVICE

STANDARD

ID	*Secondary Attribute*	*Specification*
7	Optical Drive	(optional) CD/DVD drive – minimum 24x speed
	USB	USB v2.0 port

EXPLANATION OF STANDARD

The functional requirement supported by this standard is to provide a media transfer mechanism to upload patch and software/firmware installation, upload diagnostic tools, and (with read/write capability) to download diagnostic information. Manufacturers are going toward a trend of using USB flash drives for sending patch/upgrade/firmware distributions. Therefore, the requirement is to provide both CD/DVD and USB flash drive capabilities. Note that security restrictions or concerns need to be addressed before using removable media.

The optical drive is optional because small form factor (e.g., blade) servers might not allow this type of component. These small form factor servers do, however, have the ability to support USB flash drives.

The type of drive specified is the typical expected minimum specification and type.

EVALUATION FACTORS

- Having an optical drive is beneficial.
- Faster speeds and/or Read-write capability is beneficial (additional security checks may be required for read/write capability).

IMPLEMENTATION GUIDANCE

- Preferred to provide removable media device for patch and software installation options; but small form factor (e.g., blade) servers might use other options.

3.1.8 POWER

STANDARD

ID	*Secondary Attribute*	*Specification*
8	Redundancy	N+1 redundancy on power supply
	Voltage	115-208V input voltage on power supplies
	Cord Type	IEC
	Power Efficiency	Until such time that there are Energy Star ratings available to most server classes, power efficiency should be consistent with industry trends.
	Power Consumption Documentation	Vendor-supplied data for power consumption of specified configuration at 100% load, "typical" active usage, and at idle.

EXPLANATION OF STANDARD

Redundancy – server must continue to operate if and when a power supply fails in order to enhance availability.

Voltage and cord type – driven by data center standards, 208V can more efficiently use power from utility

Power Efficiency – Certifications for power efficiency include: 80-plus (Dell & HP certify through them, but IBM does not), and Energy Star (still forming the standard for servers, but blade servers will be separate standard). The EPA establishes Energy Star standards – tier 2 appears to be the four-socket type of server while tier 1 appears to be the dual-socket server types. Some server manufacturers implement power efficiency management features including step-down of power consumption during low resource-consumption states.

Power Consumption Documentation – Preferred documentation is Energy Star data sheet and/or SPECpower_ssj2008 published results.

EVALUATION FACTORS

- Power management/economy and reporting/analysis features are beneficial.
- Lower power consumption rating at 100% utilization per workload capacity (in terms of SPECint_rate value) is beneficial.

IMPLEMENTATION GUIDANCE

- Cord type of IEC-C13 is preferred for most classes of rack-mount servers

3.1.9 COOLING FAN

STANDARD

ID	*Secondary Attribute*	*Specification*
9	Redundancy	Minimum N+1 redundancy

EXPLANATION OF STANDARD

Redundancy – server must continue to operate if and when a fan fails in order to enhance availability.

EVALUATION FACTORS

- Additional fans or separation of fans into multiple cooling zones is beneficial.
- Lower total decibel rating on the fans is beneficial.

IMPLEMENTATION GUIDANCE

3.1.10 OUT OF BAND MANAGEMENT

STANDARD

ID	*Secondary Attribute*	*Specification*
10	Type	Remote command-line and console-level access utilizing dedicated network interface
	Security	Secure IP-based remote management that complies with all VA security requirements
	Remote Power Control	Ability to remotely power on/off/reset server

EXPLANATION OF STANDARD

Required dedicated network interface is required outside of the other network interfaces specified in section 3.1.7.

Based on data center environment, operational needs, and minimum requirements that are typically satisfied through a wide variety of servers. Includes access for diagnostic capability.

Remote Power Control – Ability to power on/off/reset through server based management.

EVALUATION FACTORS

- Ability to remotely mount a drive or ISO is beneficial. This can be used for software, firmware and OS updates using remote drive mounting, and it minimizes the need for local optical drive.
- Active Directory integration for authentication is beneficial.
- Role-based access control is beneficial.
- Accounting capability is beneficial.
- Strong password compliance is beneficial.
- Accessibility via web browser is beneficial – should support Microsoft Internet Explorer.
- Support for Secure Shell (SSH) connections is beneficial – should be FIPS 140-2 compliant.
- Ability to remotely identify server with visual indicator (e.g., indicator light) is beneficial.

IMPLEMENTATION GUIDANCE

3.1.11 SERVER HARDWARE VENDOR SUPPORT

STANDARD

ID	*Secondary Attribute*	*Specification*

11	Type	Match service level requirement for hosted applications; mission-critical level support shall include 24x7 coverage 365 days/year, less than 1 hour initial engineer-level response time, 4-hour on-site response time for emergency dispatches, and 8-hour time-to-repair
	Term Length	Minimum of three years with two additional option years
	Keep Your Hard Drive	All storage drives are kept by VA in the event of warranty/service replacement

EXPLANATION OF STANDARD

Type – the level of support matches the criticality level of the hosted services. Mission-critical level is defined in this standard, and has been shown to be accommodated by major server vendors. Other levels will have to be defined to match the business expectations of that criticality level. A non-critical example would be 8 business hours/weekday phone coverage with next-day response and repair. Access to original server manufacturer support with the option to escalate to senior technicians at VA discretion is required.

Term – some components might only allow three years of support on an initial contract, but option years must be available to extend that to the expected lifetime of the system which is nominally five years.

Keep Your Hard Drive – This is a security requirement that is now a de-facto standard in VA.

EVALUATION FACTORS

- Lower on-site response time is beneficial.
- Lower initial "engineer level" contact response time is beneficial.
- Lower time-to-repair is beneficial.
- Vendor-supplied components stocked as spares near, or at, facility is beneficial.
- Penalty clauses for vendor not meeting support contract provides benefit.
- Allowance for optional parts-delivery/customer-installation of hot swappable components, memory, and drop-in cards beneficial at the VA's option.
- Access to electronic models of equipment being proposed is beneficial. Because electronic models will be used with CAD and visual design applications for datacenter planning, three dimensional models are preferred.

IMPLEMENTATION GUIDANCE

- It is preferred to have vendor-supplied total system replacements able to ship out to the site within 24 hours via overnight shipment means.

3.1.12 HARDWARE MANAGEMENT SOFTWARE

STANDARD

ID	***Secondary Attribute***	***Specification***

12	Type	Agent or agent-less server management tool accessible through SNMP
	Automated Notification	Hardware events including exceptions, diagnostics, and failures, shall be exposed through SNMP with documented MIBs
	Shared Repository	Allows uploading all information collected by the tool into shared repository using ODBC or JBDC in a manner that preserves data fields for the diagnostic information (e.g., date/time, error code, module, description, system identification, status, system configuration) Refer to Datacenter Standards, Section xx
	Support Level	Support level of management software must match support level of server hardware

EXPLANATION OF STANDARD

The tool should allow management of the hardware through a variety of means that interface with SNMP, and must be fully documented with MIBs for all relevant data (hardware events). Integration with VA's Enterprise Exception Log Service (EELS) is a consideration and is enabled through SNMP interfacing.

Ability to access event logs through database connectivity (JDBC or ODBC) is needed for future Enterprise Management Framework tools.

Service could be provided through BIOS feature, OS-level features, or other layered features, and must provide the MIB documentation for accessing the information.

EVALUATION FACTORS

- Ability to integrate with a variety of industry enterprise-level system management tools is beneficial.
- Inclusion of an enterprise-level component with shared repository and correlation analysis features is beneficial.
- Ability to provide management of agent-less systems is beneficial.

IMPLEMENTATION GUIDANCE

- Support for SNMPv3 is preferred.
- A solution design that provides secure connections outside the physical facility is preferred.
- Automated vendor notification of failure, or pre-failure activity, to initiate vendor remediation is preferred.

3.1.13 RACK INFRASTRUCTURE COMPATIBILITY

STANDARD

ID	*Secondary Attribute*	*Specification*
13	Rack Unit Measurement	1U or 2U for rack-mount servers; blade implementations will be evaluated on an individual basis.

	Rail Type	Tool-less square-hole sliding
	Cable Management	Side-reversible for non-blade solutions

EXPLANATION OF STANDARD

Rack Unit Measurement – For this class of server, occupying 4U (or two blade slots) should not be necessary; 1 slot in a half-height blade chassis is fairly equivalent to a 1U rack-mount and 1 slot in a full-height blade chassis is fairly equivalent to a 2U rack-mount. It is expected that this class of server will be a 2U rack mount or single slot full-height blade server.

Rail Type – server must have rail adapters if needed to fit the generic non-proprietary type of rack specified in this standard.

Rack infrastructure standards must match the data center equipment rack standards.

EVALUATION FACTORS

- Lower occupied rack space (per workload) is beneficial

IMPLEMENTATION GUIDANCE

3.1.14 OPERATING SYSTEM

STANDARD

ID	*Secondary Attribute*	*Specification*
14	Type	Linux preferred based on federal requirement to utilize open source operating systems, Windows acceptable, all must VA approved operating systems and compatible with hardware.
	Version	Current target version of Linux or Windows per VA TRM
	Vendor-Installed	Depends on data center policy; default is factory-installed with default settings and most recent patch level
	Support Level / Term	Must match hardware service level; term is three years with two option years unless VA has a standing enterprise software support agreement pertinent OS.

EXPLANATION OF STANDARD

Type – Linux is preferred as the most “open” type of option, but Windows is also acceptable as an industry-wide supported option. If UNIX is required Solaris is the preferred UNIX product. Solaris is preferred due to the “open” features and high level of government security certification.

Version – Operating system build, distribution, and version must be listed in the VA TRM.

Vendor-Installed – Data center policy may require local installation of the operating system. If this is not the case, the operating system should be factory-installed and have the most recent patches applied.

Support Level / Term – The operating system support level must be consistent with the criticality of the server and match hardware service level; term is three years with two option years unless VA has a standing enterprise software support agreement pertinent OS.

EVALUATION FACTORS

- Hardware support for the latest version of OS is beneficial.
- Installation of VA-supplied image is beneficial.

IMPLEMENTATION GUIDANCE

- UNIX requires exception justification and Solaris is the preferred UNIX product.
- Preference is for the latest versions of the operating system that are certified for implementation on the VA production network and supported by the hosted applications. If earlier versions are required due to compatibility limitations, these limitations should be documented for justification.

3.1.15 PLATFORM SOFTWARE / LICENSE

STANDARD

ID	*Secondary Attribute*	*Specification*
15	Multipath	Provides automated load balancing (all paths active) and fail-over; compatible with target storage array
	Data Backup	(optional) Matches requirements of the deployment location data center and SAN backup solution
	Monitoring Tool	Matches requirements of the deployment location data center
	Layered Framework	Defined by the application owner – meet license requirements for this server
	Server Management Tools	Matches requirements of the deployment location data center
	Availability / Recovery Tools	Matches the requirements for Persistent Computing standards applied to the hosted applications/services
	Security Tools	Follows Office of Information and Technology standards including compliance with FIPS, FISMA, and VA Directives.
	Virtual Management Software	Not required

EXPLANATION OF STANDARD

Multipath – Require failover capability with all paths active; Include requisite multipath software from appropriate storage vendor to provide failover and load-balancing capabilities.

Data Backup – Although this type of server primarily uses backup tools for server image backup, there may (optionally) be a need for data backup tools. A data backup tool/agent might be required by some hosted applications that have local database storage. In addition, some applications might have data stored on the SAN and is covered by the SAN backup solution that requires an agent on the server. For these servers, enterprise data backup tools are being evaluated and might lead to an enterprise standard. In the meantime, the tool selection is driven by each data center location, and must adhere to security/encryption policies. This software is installed and configured by VA and might utilize existing/enterprise licenses.

Monitoring Tool and Server Management Tool – These are the agents required by server monitoring & management tools or other monitoring tools – if agents are required. It is expected that these tools will largely be defined by the Enterprise Management Framework. The tools must meet the VA EMF standards/tools, or be able to feed information directly into the EMF tools. Until such standards are available, these tools must meet the data center standards where the tools are being placed. This software is installed and configured by VA and might utilize existing/enterprise licenses.

Availability / Recovery Tools – These include agents or other software for image backup and restoration tools, service replication tools, and fail-over tools. These will be defined in the Persistent Computing standards and by the hosted application's model for persistent computing. This software might be installed/configured through a service contract associated with the server purchase, or might be installed/configured by VA.

Security Tools – must follow Office of Information and Technology standards including compliance with FIPS, FISMA, and VA Directives, and be implemented in a way that meets Certification and Accreditation (C&A) requirements. VA Enterprise Security Solution Service (ESSS) can provide a review of proposed security tools and can be reached by contacting the Enterprise Solutions Security Service under OI&T Field Security Operations.

EVALUATION FACTORS

- Multipath – compatibility with a wide range of storage arrays / manufacturers is a benefit.

IMPLEMENTATION GUIDANCE

- This section is designed to provide guidance on what is typically purchased\installed on server platforms and should be used as a reference when making purchase\deployment decisions.

3.1.16 DATACENTER MANAGEMENT

STANDARD

ID	*Secondary Attribute*	*Specification*
16	Remote Keyboard/Video/Mouse	Ability to remotely connect to the server's external keyboard, video, and mouse ports through a networked KVM switch.

	Control	
	Remote Power Disconnect	Ability to remotely disconnect/reconnect server to power source

EXPLANATION OF STANDARD

Remote Keyboard/Video/Mouse Control – A dongle to attach the server to a networked KVM switch is required. Refer to the Datacenter Standard, Section xx

Remote Power Disconnect – Ability to remotely disconnect/reconnect power source to server through the control of individual receptacles. Refer to Datacenter Standards, Section xx.

EVALUATION FACTORS

- Ability to remotely update device firmware is beneficial.
- Accessibility via web browser is beneficial – should support Microsoft Internet Explorer.

IMPLEMENTATION GUIDANCE

3.2 TYPICAL WORKLOAD **(CLASS B)**

3.2.1 PROCESSOR

STANDARD

ID	***Secondary Attribute***	***Specification***
1	Type	x86_64 (e.g., Intel Xeon, AMD Opteron)
	Socket Count	Minimum required to meet Performance Benchmark standard
	Core Count	Amount required to meet Performance Benchmark standard
	Cache	Minimum total 2MB L2 Cache on chip (6MB or greater L3 Cache required for transactional servers)
	Virtualization Aware	Not Required

EXPLANATION OF STANDARD

Type – open/industry standard components.

Socket Count – minimize due to licensing costs per socket, have minimum count for performance/capacity.

Core Count – minimum amount expected for performance/capacity level is quad-core, more can be better. However, any amount that meets the Performance Benchmark specification is sufficient.

Cache – minimum expected for class of processor; low-power options might use lower amounts of cache.

EVALUATION FACTORS

- Currency of processor – most recent family and speed of processor is beneficial.
- Higher amounts of cache and cores are beneficial.
- Open processor slots or ability to enable additional processors for future expansion could be beneficial in larger configurations.

IMPLEMENTATION GUIDANCE

- It is recommended to have a Socket Count of two in order to minimize licensing costs that are calculated on a per socket basis. Exceptions for higher socket count should have justification documented and should consider both performance and cost.

3.2.2 PERFORMANCE

STANDARD

ID	*Secondary Attribute*	*Specification*
2	Performance Benchmark	SPEC CPU2006 Integer Rate Base value >= 200

EXPLANATION OF STANDARD

This SPEC_Int rate is represented by a large number of server platforms. Performance requirement is largely dependent on the application's workload and compute-resource consumption profile. The architecture of this server tier may provide sufficient horizontal scaling to accommodate larger workloads, especially with minimum server counts required for redundancy and limit of single-server impact. This specification is a starting point for a "typical" medium-workload application-server environment based upon review of existing tested systems in the VA environment and their capabilities represented with newer server offerings. This specification should be adjusted to meet the known workload of target hosted applications and the target architecture for its server tier.

EVALUATION FACTORS

- Higher benchmark rating value is beneficial.

IMPLEMENTATION GUIDANCE

3.2.3 MEMORY

STANDARD

ID	*Secondary Attribute*	*Specification*
3	Type	Compatible with CPU with full error correction capability including ECC; all memory modules must be of the same type and size.
	Total Memory	Minimum 48GB usable
	Speed	Minimum - match operational CPU and bus speed required to meet the Performance Benchmark requirement

EXPLANATION OF STANDARD

Type – Memory has to be highly reliable – ECC is the expected feature to allow higher levels of reliability. RDIMM is more expandable/scalable and provides additional reliability features (e.g., more error-checking features) while consuming more power (about one watt per chip) and costing more. In general, RDIMM is recommended with high memory requirements and higher availability requirements.

Total Memory – Memory requirements depend on application needs. The standard must meet the minimum requirement for memory intensive applications that could be run on this class of server. Using WebLogic as an example, minimum expected is 2GB-4GB per CPU core (as per Java configuration guidelines). With current VA application designs each JVM requires at least 4GB of memory; more if 64 bit JVMs are used. For a typical dual-socket quad-core server running WebLogic, the minimum expected requirement is therefore 48GB which would allow up to 12 JVMs. Usable memory value indicates the possibility of RAID or mirrored memory for additional redundancy. When using memory configurations above the minimum, internal storage capacity must be adjusted to support page/swap and dump requirements. This should be calculated using the formula: (2XRAM=page/swap size) + (1XRAM=DUMP) + (20GB for OS)

Speed **–** At least one operational speed of the memory must match the highest speed capability of the information bus used to retrieve memory for the CPU in order to maximize the investment into the CPU. Features can allow for step-down of speed correlated with workload consumption (power-saving feature).

EVALUATION FACTORS

- Additional memory is beneficial
- Higher density chips while avoiding a decrease in memory bus speed is preferred.
- Speed of memory takes precedence over density as long as capacity is met

IMPLEMENTATION GUIDANCE

- RDIMM is recommended for high memory requirements and higher availability requirements.

3.2.4 INTERNAL STORAGE

STANDARD

ID	*Secondary Attribute*	*Specification*
4	Type	SAS RAID controller supporting SAS drives and RAID/hot-spare features specified for this server
	Quantity of Drives	Minimum three
	RAID Configuration	Minimally require 2 disks in RAID1, plus 1 as hot spare, for page/swap, dump file, and options operating system functional role. Additional local storage needs (e.g., "temp file" storage for a transaction server or standalone storage configurations) are satisfied by additional drives in RAID level 1, 5, or 10 configuration, with a hot spare which can be shared with the page/swap disk's hot spare.
	Drive Speed	Minimum 10K RPM
	Drive Capacity	Minimum 300GB usable storage capacity
	Disk Subsystem Performance	The disk raid group must meet or exceed the performance levels specified for tests "Transaction Rate" and "Throughput Rate" Validation will be incorporated in user acceptance phase.
	Controller Cache	Minimum 256MB
	Quantity of Controllers	Minimum of one

EXPLANATION OF STANDARD

Type – SAS is preferred over SATA for the following reasons: 1) SAS uses the SCSI command set which provides better error reporting and recovery techniques than does SATA, 2) SATA throughput can be degraded on SAS adapters that use the Serial Tunneling Protocol (STP) to communicate with SATA targets, and 3) because they are engineered for continuous duty cycles and longer MTBFs, SAS drives are an accepted industry standard for enterprise server solutions. Solid state drives are beneficial due to power, cooling, performance, and reliability, but are expensive and currently are only used sparingly in external storage systems.

Quantity & RAID Configuration – requires performance for a page/swap disk and dump file, with possible additional uses; ability to continue with a drive failure (e.g., RAID1) and restore to this state with an immediately available drive (hot spare). For high performance functions (including page/swap disk), RAID1 is preferred over parity protection due to overhead of calculating parity. Although additional drives may help spread I/O across more drives for enhanced performance, additional drives also add to the cost, maintenance, and power consumption on the server. Therefore, additional drives should only be considered when they are required for performance or capacity reasons. The optimal configuration uses the fewest quantity of drives to meet the performance and reliability requirements. Additional drives are needed for additional storage needs such as

transaction server “temp file” storage or a standalone storage configuration which is provided on a separate disk per industry best practice guidance.

Drive Speed – Minimum RPM for this class of server & storage is 10K.

Drive Capacity –300GB minimum usable storage capacity is sufficient to support an operating system and/or swap and log files. This amount represents the expected minimum required for page/swap file, dump file, and optional system disk storage needs. When using memory configurations above the minimum, internal storage capacity must be adjusted to support page/swap and dump requirements. This should be calculated using the formula: (2XRAM=page/swap size) + (1XRAM=DUMP) + (20GB for OS)

Disk Subsystem Performance – The disk raid group must meet or exceed the performance levels specified below for tests “Transaction Rate” and “Throughput Rate” as described in APPENDIX G.

- Transaction Rate = 500+ IOPS
- Throughput Rate = 50+ MBPS

Controller Cache – minimum expected value for typical controllers for this class of server.

Quantity of Controllers – limited value to have more than one with only one RAID configuration; might be beneficial for more than two RAID sets if performance requires it; current technology offerings limit the ability to fail services over from one controller to another without manual intervention – future technology offerings such as dual-domain active/active configurations in SAS controllers & drives provide this capability

EVALUATION FACTORS

- Higher storage capacity is beneficial
- Higher RPM is beneficial for drive platters
- The optimal configuration uses the fewest quantity of drives to meet the performance and reliability requirements.
- Solid-state drives would be beneficial
- Higher throughput capability is beneficial
- Higher IOPS capability is beneficial
- More controllers can be beneficial
- More cache is beneficial
- Battery backup of cache is beneficial
- Support for RAID levels 0, 1, 5, 6, 0+1, and/or 1+0 is beneficial
- Separate hot spares for separate RAID groups is beneficial.

IMPLEMENTATION GUIDANCE

- SAS interface type is recommended. Battery backup of cache is recommended.

3.2.5 EXTERNAL STORAGE INTERFACE

STANDARD

ID	*Secondary Attribute*	*Specification*
5	Type	Fiber-channel with duplex LC connectors, CPU bus type is PCI-E Or Ethernet matching connector type of storage-dedicated network in data center. TOE / TCP checksum offload required.
	Port Count	Minimum of two
	Bandwidth	Minimum 4Gbps for optical or 1Gbps for copper; must meet or exceed SAN fabric requirements
	Compatibility	Support storage arrays at deployment location; supported by operating system and multipath drivers specified for this server
	SAN Boot Capability	BIOS support for SAN boot
	Blade Chassis Component (if applicable)	Dual embedded switches
	Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the SAN. Minimum two adapters with minimum one port per adapter. Support for link aggregation.

EXPLANATION OF STANDARD

- Type – fiber channel and Ethernet are the vast majority of investment for connections used in storage-dedicated networks in VA data centers, and exceptions should be justified including impact to existing infrastructure investment. Typical Ethernet connector types are RJ-45 for copper and LC for optical. When Ethernet is used, TOE / TCP checksum offload is required. Drivers are more readily available for Windows than Linux servers.
- Port Count – minimum required for port-level redundancy / load-balancing
- Bandwidth – 4Gbps is standard for current investment of optical fabrics in the storage systems and infrastructure in VA data centers and matches the minimum current industry standard; similarly 1Gbps is standard for current investments and minimum current industry practices when using copper-based storage fabrics.
- Compatibility – Support storage subsystems at deployment locations. Check adapter compatibility for storage at deployment location.
- Blade Chassis Component – A primary benefit of blade servers is their ability to reduce datacenter cabling. Dual fiber-channel switches (or Ethernet if applicable) are required. Pass-thru configurations require an exception. Each blade chassis shall have redundant sets of embedded switches to provide independent SAN

connections to each blade. Each embedded switch set shall have the ability to configure redundant and/or trunked ISL/uplink ports. Minimum number of uplink ports in chassis is two times the number of embedded switches. Refer to Storage Standards, Section xx.

- Hardware Redundancy – Fail-over and load-balancing are functions of the multipath tool which is covered in section 3.2.16

EVALUATION FACTORS

- Redundancy across physically independent interfaces is beneficial
- Support for 8Gbps optical or 10Gbps copper bandwidth is beneficial

IMPLEMENTATION GUIDANCE

- For a server with high criticality of SAN storage connectivity, redundant adapter cards are recommended.
- Interface for iSCSI can be termed an HBA interface or a NIC interface. Where higher performance is a concern, the HBA interface is recommended.
- Bandwidth spread across multiple ports on independent cards may be required for performance reasons

3.2.6 SERVER COMMUNICATIONS INTERFACE

STANDARD

ID	*Secondary Attribute*	*Specification*
6	Type	Ethernet matching connector type for network in data center
	Port Count	Minimum of 2 ports; at least 2 more if private cluster interconnect is required
	Bandwidth	Minimum 1Gbps
	Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the LAN; supports redundant ports (may include embedded ports) for each interface function (e.g., public network function and cluster interconnect function)
	Blade Chassis Component (if applicable)	Dual embedded switches

EXPLANATION OF STANDARD

Type – Ethernet is the vast majority of investment for connections used in data center networks in VA, and exceptions should be justified including impact to existing infrastructure investment. Typical Ethernet connector types are RJ-45 for copper and LC for optical with a future trend towards copper only.

Port Count – base set of 2 ports are for general “public” Ethernet; other specific interfaces may be required depending on the profile of the workload/application. Cluster interconnects may need 2 or more ports, depending

on the amount of cluster traffic (e.g., Oracle clusters send data over the cluster interconnect, simple Microsoft clustering only send "heartbeat" traffic).

Bandwidth – expected standard interface on network switches is 1Gbps. There is a potential for 10Gbps Ethernet.

Hardware Redundancy – embedded ports count as one "card" which is sufficient for this class of server except where additional ports are required. If additional ports are required, all interface uses are spread across two "cards" (including at least one add-on NIC). The goal is physical separation of the chipset/ASIC used on the redundant port, for each function. Support for "simple failover" (with a single IP address) and independent usage (with multiple IP addresses) is required.

Blade Chassis Component – Where applicable, blade-server chassis should be configured with two switches to limit port and cabling needs external to the server, unless there is a specific workload/application requirement for high throughput that requires pass-through capability.

EVALUATION FACTORS

- Additional cards can be beneficial
- Support for TCP Offload Engine is beneficial
- Support for active/active load balancing is beneficial

IMPLEMENTATION GUIDANCE

- One or two additional ports are recommended for network backups if the impact of backup traffic may affect other services on the "public" network. One port is sufficient if only one port is active and uses another port for fail-over.
- TCP Offload Engine (TOE) is preferred for all servers. Drivers are more readily available for Windows than Linux servers.
- If there is a specific workload/application requirement for high throughput that requires pass-through capability on a blade-server chassis, then this option should be configured, but the chassis switch is preferred otherwise.
- Bandwidth spread across multiple ports on independent cards may be required for performance reasons

3.2.7 REMOVABLE MEDIA DEVICE

STANDARD

ID	*Secondary Attribute*	*Specification*
7	Optical Drive	(optional) CD/DVD drive – minimum 24x speed
	USB	USB v2.0 port

EXPLANATION OF STANDARD

The functional requirement supported by this standard is to provide a media transfer mechanism to upload patch and software/firmware installation, upload diagnostic tools, and (with read/write capability) to download

diagnostic information. Manufacturers are going toward a trend of using USB flash drives for sending patch/upgrade/firmware distributions. Therefore, the requirement is to provide both CD/DVD and USB flash drive capabilities. Note that security restrictions or concerns need to be addressed before using removable media.

The optical drive is optional because small form factor (e.g., blade) servers might not allow this type of component. These small form factor servers do, however, have the ability to support USB flash drives.

The type of drive specified is the typical expected minimum specification and type.

EVALUATION FACTORS

- Having an optical drive is beneficial.
- Faster speeds and/or Read-write capability is beneficial (additional security checks may be required for read/write capability).

IMPLEMENTATION GUIDANCE

- Preferred to provide removable media device for patch and software installation options; but small form factor (e.g., blade) servers might use other options.

3.2.8 POWER

STANDARD

ID	*Secondary Attribute*	*Specification*
8	Redundancy	N+1 redundancy on power supply
	Voltage	115-208V input voltage on power supplies
	Cord Type	IEC
	Power Efficiency	Until such time that there are Energy Star ratings available to most server classes, power efficiency should be consistent with industry trends.
	Power Consumption Documentation	Vendor-supplied data for power consumption of specified configuration at 100% load, “typical” active usage, and at idle.

EXPLANATION OF STANDARD

Redundancy – server must continue to operate if and when a power supply fails in order to enhance availability.

Voltage and cord type – driven by data center standards, 208V can more efficiently use power from utility

Power Efficiency – Certifications for power efficiency include: 80-plus (Dell & HP certify through them, but IBM does not), and Energy Star (still forming the standard for servers, but blade servers will be separate standard). The EPA establishes Energy Star standards – tier 2 appears to be the four-socket type of server while tier 1 appears to

be the dual-socket server types. Some server manufacturers implement power efficiency management features including step-down of power consumption during low resource-consumption states.

Power Consumption Documentation – Preferred documentation is Energy Star data sheet and/or SPECpower_ssj2008 published results.

EVALUATION FACTORS

- Power management/economy and reporting/analysis features are beneficial.
- Lower power consumption rating at 100% utilization per workload capacity (in terms of SPECint_rate value) is beneficial.

IMPLEMENTATION GUIDANCE

- Cord type of IEC-C13 is preferred for most classes of rack-mount servers

3.2.9 COOLING FAN

STANDARD

ID	*Secondary Attribute*	*Specification*
9	Redundancy	Minimum N+1 redundancy

EXPLANATION OF STANDARD

- Redundancy – server must continue to operate if and when a fan fails in order to enhance availability.

EVALUATION FACTORS

- Additional fans or separation of fans into multiple cooling zones is beneficial.
- Lower total decibel rating on the fans is beneficial.

IMPLEMENTATION GUIDANCE

3.2.10 OUT OF BAND MANAGEMENT

STANDARD

ID	*Secondary Attribute*	*Specification*
10	Type	Remote command-line and console-level access utilizing dedicated network interface
	Security	Secure IP-based remote management that complies with all VA security requirements
	Remote Power Control	Ability to remotely power on/off/reset server

EXPLANATION OF STANDARD

Required dedicated network interface is required outside of the other network interfaces specified in section 3.2.7.

Based on data center environment, operational needs, and minimum requirements that are typically satisfied through a wide variety of servers. Includes access for diagnostic capability.

Remote Power Control – Ability to power on/off/reset through server based management.

EVALUATION FACTORS

- Ability to remotely mount a drive or ISO is beneficial. This can be used for software, firmware and OS updates using remote drive mounting, and it minimizes the need for local optical drive.
- Active Directory integration for authentication is beneficial.
- Role-based access control is beneficial.
- Accounting capability is beneficial.
- Strong password compliance is beneficial.
- Accessibility via web browser is beneficial – should support Microsoft Internet Explorer.
- Support for Secure Shell (SSH) connections is beneficial – should be FIPS 140-2 compliant.
- Ability to remotely identify server with visual indicator (e.g., indicator light) is beneficial.

IMPLEMENTATION GUIDANCE

3.2.11 SERVER HARDWARE VENDOR SUPPORT

STANDARD

ID	*Secondary Attribute*	*Specification*
11	Type	Match service level requirement for hosted applications; mission-critical level

		support shall include 24x7 coverage 365 days/year, less than 1 hour initial engineer-level response time, 4-hour on-site response time for emergency dispatches, and 8-hour time-to-repair
	Term Length	Minimum of three years with two additional option years
	Keep Your Hard Drive	All storage drives are kept by VA in the event of warranty/service replacement

EXPLANATION OF STANDARD

Type – the level of support matches the criticality level of the hosted services. Mission-critical level is defined in this standard, and has been shown to be accommodated by major server vendors. Other levels will have to be defined to match the business expectations of that criticality level. A non-critical example would be 8 business hours/weekday phone coverage with next-day response and repair. Access to original server manufacturer support with the option to escalate to senior technicians at VA discretion is required.

Term – some components might only allow three years of support on an initial contract, but option years must be available to extend that to the expected lifetime of the system which is nominally five years.

Keep Your Hard Drive – This is a security requirement that is now a de-facto standard in VA.

EVALUATION FACTORS

- Lower on-site response time is beneficial.
- Lower initial "engineer level" contact response time is beneficial.
- Lower time-to-repair is beneficial.
- Vendor-supplied components stocked as spares near, or at, facility is beneficial.
- Penalty clauses for vendor not meeting support contract provides benefit.
- Allowance for optional parts-delivery/customer-installation of hot swappable components, memory, and drop-in cards beneficial at the VA's option.
- Access to electronic models of equipment being proposed is beneficial. Because electronic models will be used with CAD and visual design applications for datacenter planning, three dimensional models are preferred.

IMPLEMENTATION GUIDANCE

- It is preferred to have vendor-supplied total system replacements able to ship out to the site within 24 hours via overnight shipment means.

3.2.12 HARDWARE MANAGEMENT SOFTWARE

STANDARD

ID	***Secondary Attribute***	***Specification***
12	Type	Agent or agent-less server management tool accessible through SNMP

	Automated Notification	Hardware events including exceptions, diagnostics, and failures, shall be exposed through SNMP with documented MIBs
	Shared Repository	Allows uploading all information collected by the tool into shared repository using ODBC or JBDC in a manner that preserves data fields for the diagnostic information (e.g., date/time, error code, module, description, system identification, status, system configuration)
	Support Level	Support level of management software must match support level of server hardware

EXPLANATION OF STANDARD

The tool should allow management of the hardware through a variety of means that interface with SNMP, and must be fully documented with MIBs for all relevant data (hardware events). Integration with VA's Enterprise Exception Log Service (EELS) is a consideration and is enabled through SNMP interfacing.

Ability to access event logs through database connectivity (JDBC or ODBC) is needed for future Enterprise Management Framework tools.

Service could be provided through BIOS feature, OS-level features, or other layered features, and must provide the MIB documentation for accessing the information.

EVALUATION FACTORS

- Ability to integrate with a variety of industry enterprise-level system management tools is beneficial.
- Inclusion of an enterprise-level component with shared repository and correlation analysis features is beneficial.
- Ability to provide management of agent-less systems is beneficial.

IMPLEMENTATION GUIDANCE

- Support for SNMPv3 is preferred.
- A solution design that provides secure connections outside the physical facility is preferred.
- Automated vendor notification of failure, or pre-failure activity, to initiate vendor remediation is preferred.

3.2.13 RACK INFRASTRUCTURE COMPATIBILITY

STANDARD

ID	*Secondary Attribute*	*Specification*
13	Rack Unit Measurement	1U or 2U for rack-mount servers; blade implementations will be evaluated on an individual basis.
	Rail Type	Tool-less square-hole sliding

	Cable Management	Side-reversible for non-blade solutions

EXPLANATION OF STANDARD

Rack Unit Measurement – For this class of server, occupying 4U (or two blade slots) should not be necessary; 1 slot in a half-height blade chassis is fairly equivalent to a 1U rack-mount and 1 slot in a full-height blade chassis is fairly equivalent to a 2U rack-mount. It is expected that this class of server will be a 2U rack mount or single slot full-height blade server.

Rail Type – server must have rail adapters if needed to fit the generic non-proprietary type of rack specified in this standard.

Rack infrastructure standards must match the data center equipment rack standards.

EVALUATION FACTORS

- Lower occupied rack space (per workload) is beneficial

IMPLEMENTATION GUIDANCE

3.2.14 OPERATING SYSTEM

STANDARD

ID	*Secondary Attribute*	*Specification*
14	Type	Linux preferred based on federal requirement to utilize open source operating systems, Windows acceptable, all must VA approved operating systems and compatible with hardware.
	Version	Current target version of Linux or Windows per VA TRM
	Vendor-Installed	Depends on data center policy; default is factory-installed with default settings and most recent patch level
	Support Level / Term	Must match hardware service level; term is three years with two option years unless VA has a standing enterprise software support agreement pertinent OS.

EXPLANATION OF STANDARD

Type – Linux is preferred as the most “open” type of option, but Windows is also acceptable as an industry-wide supported option. If UNIX is required Solaris is the preferred UNIX product. Solaris is preferred due to the “open” features and high level of government security certification.

Version – Operating system build, distribution, and version must be listed in the VA TRM.

Vendor-Installed – Data center policy may require local installation of the operating system. If this is not the case, the operating system should be factory-installed and have the most recent patches applied.

Support Level / Term – The operating system support level must be consistent with the criticality of the server and match hardware service level; term is three years with two option years unless VA has a standing enterprise software support agreement pertinent OS.

EVALUATION FACTORS

- Hardware support for the latest version of OS is beneficial.
- Installation of VA-supplied image is beneficial.

IMPLEMENTATION GUIDANCE

- UNIX requires exception justification and Solaris is the preferred UNIX product.
- Preference is for the latest versions of the operating system that are certified for implementation on the VA production network and supported by the hosted applications. If earlier versions are required due to compatibility limitations, these limitations should be documented for justification.

3.2.15 PLATFORM SOFTWARE / LICENSE

STANDARD

ID	*Secondary Attribute*	*Specification*
15	Multipath	Provides automated load balancing (all paths active) and fail-over; compatible with target storage array
	Data Backup	(optional) Matches requirements of the deployment location data center and SAN backup solution
	Monitoring Tool	Matches requirements of the deployment location data center
	Layered Framework	Defined by the application owner – meet license requirements for this server
	Server Management Tools	Matches requirements of the deployment location data center
	Availability / Recovery Tools	Matches the requirements for Persistent Computing standards applied to the hosted applications/services
	Security Tools	Follows Office of Information and Technology standards including compliance with FIPS, FISMA, and VA Directives.
	Virtual Management Software	Not Required

EXPLANATION OF STANDARD

Multipath – Require failover capability with all paths active; Include requisite multipath software from appropriate storage vendor to provide failover and load-balancing capabilities.

Data Backup – Although this type of server primarily uses backup tools for server image backup, there may (optionally) be a need for data backup tools. A data backup tool/agent might be required by some hosted applications that have local database storage. In addition, some applications might have data stored on the SAN and is covered by the SAN backup solution that requires an agent on the server. For these servers, enterprise data backup tools are being evaluated and might lead to an enterprise standard. In the meantime, the tool selection is driven by each data center location, and must adhere to security/encryption policies. This software is installed and configured by VA and might utilize existing/enterprise licenses.

Monitoring Tool and Server Management Tool – These are the agents required by server monitoring & management tools or other monitoring tools – if agents are required. It is expected that these tools will largely be defined by the Enterprise Management Framework. The tools must meet the VA EMF standards/tools, or be able to feed information directly into the EMF tools. Until such standards are available, these tools must meet the data center standards where the tools are being placed. This software is installed and configured by VA and might utilize existing/enterprise licenses.

Availability / Recovery Tools – These include agents or other software for image backup and restoration tools, service replication tools, and fail-over tools. These will be defined in the Persistent Computing standards and by the hosted application's model for persistent computing. This software might be installed/configured through a service contract associated with the server purchase, or might be installed/configured by VA.

Security Tools – must follow Office of Information and Technology standards including compliance with FIPS, FISMA, and VA Directives, and be implemented in a way that meets Certification and Accreditation (C&A) requirements. VA Enterprise Security Solution Service (ESSS) can provide a review of proposed security tools and can be reached by contacting the Enterprise Solutions Security Service under OI&T Field Security Operations.

EVALUATION FACTORS

- Multipath – compatibility with a wide range of storage arrays / manufacturers is a benefit.

IMPLEMENTATION GUIDANCE

- .This section is designed to provide guidance on what is typically purchased\installed on server platforms and should be used as a reference when making purchase\deployment decisions.

3.2.16 DATACENTER MANAGEMENT

STANDARD

ID	*Secondary Attribute*	*Specification*
16	Remote Keyboard/Video/Mouse	Ability to remotely connect to the server's external keyboard, video, and mouse ports through a networked KVM switch.

	Control	
	Remote Power Disconnect	Ability to remotely disconnect/reconnect server to power source

EXPLANATION OF STANDARD

Remote Keyboard/Video/Mouse Control – A dongle to attach the server to a networked KVM switch is required. Refer to the Datacenter Standard, Section xx

Remote Power Disconnect – Ability to remotely disconnect/reconnect power source to server through the control of individual receptacles. Refer to Datacenter Standards, Section xx.

EVALUATION FACTORS

- Ability to remotely update device firmware is beneficial.
- Accessibility via web browser is beneficial – should support Microsoft Internet Explorer.

IMPLEMENTATION GUIDANCE

3.3 HEAVY WORKLOAD (**CLASS A**)

3.3.1 PROCESSOR

STANDARD

ID	***Secondary Attribute***	***Specification***
1	Type	x86_64 (e.g., Intel Xeon, AMD Opteron)
	Socket Count	Minimum required to meet Performance Benchmark standard
	Core Count	Amount required to meet Performance Benchmark standard
	Cache	Minimum total 2MB L2 Cache on chip (6MB or greater L3 Cache required for transactional servers)
	Virtualization Aware	Required to facilitate live migration

EXPLANATION OF STANDARD

Type – open/industry standard components.

Socket Count – minimize due to licensing costs per socket, have minimum count for performance/capacity.

Core Count – minimum amount expected for performance/capacity level is quad-core, more can be better. However, any amount that meets the Performance Benchmark specification is sufficient.

Cache – minimum expected for class of processor; low-power options might use lower amounts of cache.

Virtualization Aware – CPU manufacturers offer technology that allows easy migration of running virtual guests from one physical host to another. Servers in this class must provide that functionality.

EVALUATION FACTORS

- Currency of processor – most recent family and speed of processor is beneficial.
- Higher amounts of cache and cores are beneficial.
- Open processor slots or ability to enable additional processors for future expansion could be beneficial in larger configurations.

IMPLEMENTATION GUIDANCE

- It is recommended to have a Socket Count of four or less in order to minimize licensing costs that are calculated on a per socket basis. Exceptions for higher socket count should have justification documented and should consider both performance and cost.

3.3.2 PERFORMANCE

STANDARD

ID	*Secondary Attribute*	*Specification*
2	Performance Benchmark	SPEC CPU2006 Integer Rate Base value >= 250

EXPLANATION OF STANDARD

This SPEC_Int rate is represented by a large number of server platforms. Performance requirement is largely dependent on the application's workload and compute-resource consumption profile. This specification is a starting point for a "heavy duty" high-workload transactional or virtual server environment based upon review of existing tested systems in the VA environment and their capabilities represented with newer server offerings. This specification should be adjusted to meet the known workload of target hosted applications and the target architecture for its server tier.

EVALUATION FACTORS

- Higher benchmark rating value is beneficial.

IMPLEMENTATION GUIDANCE

3.3.3 MEMORY

STANDARD

ID	*Secondary Attribute*	*Specification*
3	Type	Compatible with CPU with full error correction capability including ECC; all memory modules must be of the same type and size.
	Total Memory	Minimum 128GB usable
	Speed	Minimum - match operational CPU and bus speed required to meet the Performance Benchmark requirement

EXPLANATION OF STANDARD

Type – Memory has to be highly reliable – ECC is the expected feature to allow higher levels of reliability. RDIMM is more expandable/scalable and provides additional reliability features (e.g., more error-checking features) while consuming more power (about one watt per chip) and costing more. In general, RDIMM is recommended with high memory requirements and higher availability requirements.

Total Memory – Memory requirements depend on application needs. The standard must meet the minimum requirement for memory intensive applications that could be run on this class of server. Using virtual hosting as an example, the minimum standard would allow for 4 to 8 virtual guests configured with 16GB to 32GB each. Usable memory value indicates the possibility of RAID or mirrored memory for additional redundancy. This class of server should provide at least 8 environments in a medium application environment. When using memory configurations above the minimum, internal storage capacity must be adjusted to support page/swap and dump requirements. This should be calculated using the formula: (2XRAM=page/swap size) + (1XRAM=DUMP) + (20GB for OS)

Speed – At least one operational speed of the memory must match the highest speed capability of the information bus used to retrieve memory for the CPU in order to maximize the investment into the CPU. Features can allow for step-down of speed correlated with workload consumption (power-saving feature).

EVALUATION FACTORS

- Additional memory is beneficial
- Higher density chips while avoiding a decrease in memory bus speed is preferred.
- Speed of memory takes precedence over density as long as capacity is met

IMPLEMENTATION GUIDANCE

- RDIMM is recommended for high memory requirements and higher availability requirements.

3.3.4 INTERNAL STORAGE

STANDARD

ID	*Secondary Attribute*	*Specification*
4	Type	SAS RAID controller supporting SAS drives and RAID/hot-spare features specified for this server
	Quantity of Drives	Minimum three
	RAID Configuration	Minimally require 2 disks in RAID1, plus 1 as hot spare, for page/swap, dump file, and options operating system functional role. Additional local storage needs (e.g., "temp file" storage for a transaction server or standalone storage configurations) are satisfied by additional drives in RAID level 1, 5, or 10 configuration, with a hot spare which can be shared with the page/swap disk's hot spare. When this class of server is configured as a transaction server, a separate RAID controller or second channel on the same controller is required for the "temp file" RAID group
	Drive Speed	Minimum 10K RPM
	Drive Capacity	Minimum 450GB usable storage capacity
	Disk Subsystem Performance	The disk raid group must meet or exceed the performance levels specified for tests "Transaction Rate" and "Throughput Rate" Validation will be incorporated in user acceptance phase.
	Controller Cache	Minimum 256MB
	Quantity of Controllers	Minimum of one

EXPLANATION OF STANDARD

Type – SAS is preferred over SATA for the following reasons: 1) SAS uses the SCSI command set which provides better error reporting and recovery techniques than does SATA, 2) SATA throughput can be degraded on SAS adapters that use the Serial Tunneling Protocol (STP) to communicate with SATA targets, and 3) because they are engineered for continuous duty cycles and longer MTBFs, SAS drives are an accepted industry standard for enterprise server solutions. Solid state drives are beneficial due to power, cooling, performance, and reliability, but are expensive and currently are only used sparingly in external storage systems.

Quantity & RAID Configuration – requires performance for a page/swap disk and dump file, with possible additional uses; ability to continue with a drive failure (e.g., RAID1) and restore to this state with an immediately available drive (hot spare). For high performance functions (including page/swap disk), RAID1 is preferred over parity protection due to overhead of calculating parity. Although additional drives may help spread I/O across more drives for enhanced performance, additional drives also add to the cost, maintenance, and power consumption on the server. Therefore, additional drives should only be considered when they are required for performance or capacity reasons. The optimal configuration uses the fewest quantity of drives to meet the

performance and reliability requirements. Additional drives are needed for additional storage needs such as transaction server "temp file" storage or a standalone storage configuration which is provided on a separate disk per industry best practice guidance.

Drive Speed – Minimum RPM for this class of server & storage is 10K.

Drive Capacity – 450GB minimum usable storage capacity is sufficient to support an operating system and/or swap and log files. This amount represents the expected minimum required for page/swap file, dump file, and optional system disk storage needs. When using memory configurations above the minimum, internal storage capacity must be adjusted to support page/swap and dump requirements. This should be calculated using the formula: (2XRAM=page/swap size) + (1XRAM=DUMP) + (20GB for OS)

Disk Subsystem Performance – The disk raid group must meet or exceed the performance levels specified below for tests "Transaction Rate" and "Throughput Rate" as described in APPENDIX G.

- Transaction Rate = 500+ IOPS
- Throughput Rate = 50+ MBPS

Controller Cache – minimum expected value for typical controllers for this class of server.

Quantity of Controllers – limited value to have more than one with only one RAID configuration; might be beneficial for more than two RAID sets if performance requires it; current technology offerings limit the ability to fail services over from one controller to another without manual intervention – future technology offerings such as dual-domain active/active configurations in SAS controllers & drives provide this capability

EVALUATION FACTORS

- Higher storage capacity is beneficial
- Higher RPM is beneficial for drive platters
- The optimal configuration uses the fewest quantity of drives to meet the performance and reliability requirements.
- Solid-state drives would be beneficial
- Higher throughput capability is beneficial
- Higher IOPS capability is beneficial
- More controllers can be beneficial
- More cache is beneficial
- Battery backup of cache is beneficial
- Support for RAID levels 0, 1, 5, 6, 0+1, and/or 1+0 is beneficial
- Separate hot spares for separate RAID groups is beneficial.

IMPLEMENTATION GUIDANCE

- SAS interface type is recommended. Battery backup of cache is recommended.

3.3.5 EXTERNAL STORAGE INTERFACE

STANDARD

ID	*Secondary Attribute*	*Specification*
5	Type	Fiber-channel with duplex LC connectors, CPU bus type is PCI-E Or Ethernet matching connector type of storage-dedicated network in data center. TOE / TCP checksum offload required.
	Port Count	Minimum of 2 independent cards with a minimum of 2 ports each
	Bandwidth	Minimum 4Gbps for optical or 1Gbps for copper; must meet or exceed SAN fabric requirements
	Compatibility	Support storage arrays at deployment location; supported by operating system and multipath drivers specified for this server
	SAN Boot Capability	BIOS support for SAN boot
	Blade Chassis Component (if applicable)	Dual embedded switches
	Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the SAN. Minimum two adapters with minimum one port per adapter. Support for link aggregation.

EXPLANATION OF STANDARD

- Type – fiber channel and Ethernet are the vast majority of investment for connections used in storage-dedicated networks in VA data centers, and exceptions should be justified including impact to existing infrastructure investment. Typical Ethernet connector types are RJ-45 for copper and LC for optical. When Ethernet is used, TOE / TCP checksum offload is required. Drivers are more readily available for Windows than Linux servers.
- Port Count – minimum required for port-level redundancy / load-balancing. Provides the flexibility to improve the throughput and performance without further hardware changes. Can initially configure with only 2 active. Potential benefit for separating the ports used for booting and data. Allows increased workload of multiple database servers in shared server environment.
- Bandwidth – 4Gbps is standard for current investment of optical fabrics in the storage systems and infrastructure in VA data centers and matches the minimum current industry standard; similarly 1Gbps is standard for current investments and minimum current industry practices when using copper-based storage fabrics.
- Compatibility – Support storage subsystems at deployment locations. Check adapter compatibility for storage at deployment location. Don’t want to provide a list of brand names.

- Blade Chassis Component – A primary benefit of blade servers is their ability to reduce datacenter cabling. Dual fiber-channel switches (or Ethernet if applicable) are required. Pass-thru configurations require an exception. Each blade chassis shall have redundant sets of embedded switches to provide independent SAN connections to each blade. Each embedded switch set shall have the ability to configure redundant and/or trunked ISL/uplink ports. Minimum number of uplink ports in chassis is two times the number of embedded switches. Refer to Storage Standards, Section xx.
- Hardware Redundancy – Fail-over and load-balancing are functions of the multipath tool which is covered in section 3.3.16

EVALUATION FACTORS

- Redundancy across physically independent interfaces is beneficial
- Support for 8Gbps optical or 10Gbps copper bandwidth is beneficial
- Four independent cards with a minimum of 2 ports each is preferred.

IMPLEMENTATION GUIDANCE

- For a server with high criticality of SAN storage connectivity, redundant adapter cards are recommended.
- Interface for iSCSI can be termed an HBA interface or a NIC interface. Where higher performance is a concern, the HBA interface is recommended.
- Bandwidth spread across multiple ports on independent cards may be required for performance reasons

3.3.6 SERVER COMMUNICATIONS INTERFACE

STANDARD

ID	*Secondary Attribute*	*Specification*
6	Type	Ethernet matching connector type for network in data center
	Port Count	Minimum of 6 ports; more may be necessary depending on virtual infrastructure requirements
	Bandwidth	Minimum 1Gbps
	Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the LAN; supports physically independent redundant ports including embedded ports for each interface function (e.g., public network function and cluster interconnect function)
	Blade Chassis Component (if applicable)	Dual embedded switches

EXPLANATION OF STANDARD

Type – Ethernet is the vast majority of investment for connections used in data center networks in VA, and exceptions should be justified including impact to existing infrastructure investment. Typical Ethernet connector types are RJ-45 for copper and LC for optical with a future trend towards copper only.

Port Count – when used in a virtual environment, a base set of 4 ports are used for virtual guests, 1 is for the console, and 1 for the live migration. Other specific interfaces may be required depending on the profile of the workload/application. When not used in a virtual environment, cluster interconnects may need 2 or more ports, depending on the amount of cluster traffic (e.g., Oracle clusters send data over the cluster interconnect, simple Microsoft clustering only send "heartbeat" traffic).

Bandwidth – expected standard interface on network switches is 1Gbps. There is a potential for 10Gbps Ethernet.

Hardware Redundancy – embedded ports count as one "card"; all interface uses are spread across two "cards" (including at least one add-on NIC). The goal is physical separation of the chipset/ASIC used on the redundant port, for each function. Support for "simple failover" (with a single IP address) and independent usage (with multiple IP addresses) is required.

Blade Chassis Component – Where applicable, blade-server chassis should be configured with two switches to limit port and cabling needs external to the server, unless there is a specific workload/application requirement for high throughput that requires pass-through capability.

EVALUATION FACTORS

- Additional cards can be beneficial
- Support for TCP Offload Engine is beneficial
- Support for active/active load balancing is beneficial

IMPLEMENTATION GUIDANCE

- One or two additional ports are recommended for network backups if the impact of backup traffic may affect other services on the "public" network. One port is sufficient if only one port is active and uses another port for fail-over.
- TCP Offload Engine (TOE) is preferred for all servers. Drivers are more readily available for Windows than Linux servers.
- If there is a specific workload/application requirement for high throughput that requires pass-through capability on a blade-server chassis, then this option should be configured, but the chassis switch is preferred otherwise.
- Bandwidth spread across multiple ports on independent cards may be required for performance reasons

3.3.7 REMOVABLE MEDIA DEVICE

STANDARD

ID	*Secondary Attribute*	*Specification*

7	Optical Drive	(optional) CD/DVD drive – minimum 24x speed
	USB	USB v2.0 port

EXPLANATION OF STANDARD

The functional requirement supported by this standard is to provide a media transfer mechanism to upload patch and software/firmware installation, upload diagnostic tools, and (with read/write capability) to download diagnostic information. Manufacturers are going toward a trend of using USB flash drives for sending patch/upgrade/firmware distributions. Therefore, the requirement is to provide both CD/DVD and USB flash drive capabilities. Note that security restrictions or concerns need to be addressed before using removable media.

The optical drive is optional because small form factor (e.g., blade) servers might not allow this type of component. These small form factor servers do, however, have the ability to support USB flash drives.

The type of drive specified is the typical expected minimum specification and type.

EVALUATION FACTORS

- Having an optical drive is beneficial.
- Faster speeds and/or Read-write capability is beneficial (additional security checks may be required for read/write capability).

IMPLEMENTATION GUIDANCE

- Preferred to provide removable media device for patch and software installation options; but small form factor (e.g., blade) servers might use other options.

3.3.8 POWER

STANDARD

ID	***Secondary Attribute***	***Specification***
8	Redundancy	N+1 redundancy on power supply
	Voltage	208V input voltage on power supply
	Cord Type	IEC
	Power Efficiency	Until such time that there are Energy Star ratings available to most server classes, power efficiency should be consistent with industry trends.
	Power Consumption Documentation	Vendor-supplied data for power consumption of specified configuration at 100% load, “typical” active usage, and at idle.

EXPLANATION OF STANDARD

Redundancy – server must continue to operate if and when a power supply fails in order to enhance availability.

Voltage and cord type – driven by data center standards, 208V can more efficiently use power from utility

Power Efficiency – Certifications for power efficiency include: 80-plus (Dell & HP certify through them, but IBM does not), and Energy Star (still forming the standard for servers, but blade servers will be separate standard). The EPA establishes Energy Star standards – tier 2 appears to be the four-socket type of server while tier 1 appears to be the dual-socket server types. Some server manufacturers implement power efficiency management features including step-down of power consumption during low resource-consumption states.

Power Consumption Documentation – Preferred documentation is Energy Star data sheet and/or SPECpower_ssj2008 published results.

EVALUATION FACTORS

- Power management/economy and reporting/analysis features are beneficial.
- Lower power consumption rating at 100% utilization per workload capacity (in terms of SPECint_rate value) is beneficial.

IMPLEMENTATION GUIDANCE

- Cord type of IEC-C13 is preferred for most classes of rack-mount servers

3.3.9 COOLING FAN

STANDARD

ID	*Secondary Attribute*	*Specification*
9	Redundancy	Minimum N+1 redundancy

EXPLANATION OF STANDARD

Redundancy – server must continue to operate if and when a fan fails in order to enhance availability.

EVALUATION FACTORS

- Additional fans or separation of fans into multiple cooling zones is beneficial.
- Lower total decibel rating on the fans is beneficial.

IMPLEMENTATION GUIDANCE

3.3.10 OUT OF BAND MANAGEMENT

STANDARD

ID	*Secondary Attribute*	*Specification*
10	Type	Remote command-line and console-level access utilizing dedicated network interface
	Security	Secure IP-based remote management that complies with all VA security requirements
	Remote Power Control	Ability to remotely power on/off/reset server

EXPLANATION OF STANDARD

Required dedicated network interface is required outside of the other network interfaces specified in section 3.2.7.

Based on data center environment, operational needs, and minimum requirements that are typically satisfied through a wide variety of servers. Includes access for diagnostic capability.

Remote Power Control – Ability to power on/off/reset through server based management.

EVALUATION FACTORS

- Ability to remotely mount a drive or ISO is beneficial. This can be used for software, firmware and OS updates using remote drive mounting, and it minimizes the need for local optical drive.
- Active Directory integration for authentication is beneficial.
- Role-based access control is beneficial.
- Accounting capability is beneficial.
- Strong password compliance is beneficial.
- Accessibility via web browser is beneficial – should support Microsoft Internet Explorer.
- Support for Secure Shell (SSH) connections is beneficial – should be FIPS 140-2 compliant.
- Ability to remotely identify server with visual indicator (e.g., indicator light) is beneficial.

IMPLEMENTATION GUIDANCE

3.3.11 SERVER HARDWARE VENDOR SUPPORT

STANDARD

ID	*Secondary Attribute*	*Specification*
11	Type	Match service level requirement for hosted applications; mission-critical level

		support shall include 24x7 coverage 365 days/year, less than 1 hour initial engineer-level response time, 4-hour on-site response time for emergency dispatches, and 8-hour time-to-repair
	Term Length	Minimum of three years with two additional option years
	Keep Your Hard Drive	All storage drives are kept by VA in the event of warranty/service replacement

EXPLANATION OF STANDARD

Type – the level of support matches the criticality level of the hosted services. Mission-critical level is defined in this standard, and has been shown to be accommodated by major server vendors. Other levels will have to be defined to match the business expectations of that criticality level. A non-critical example would be 8 business hours/weekday phone coverage with next-day response and repair. Access to original server manufacturer support with the option to escalate to senior technicians at VA discretion is required.

Term – some components might only allow three years of support on an initial contract, but option years must be available to extend that to the expected lifetime of the system which is nominally five years.

Keep Your Hard Drive – This is a security requirement that is now a de-facto standard in VA.

EVALUATION FACTORS

- Lower on-site response time is beneficial.
- Lower initial "engineer level" contact response time is beneficial.
- Lower time-to-repair is beneficial.
- Vendor-supplied components stocked as spares near, or at, facility is beneficial.
- Penalty clauses for vendor not meeting support contract provides benefit.
- Allowance for optional parts-delivery/customer-installation of hot swappable components, memory, and drop-in cards beneficial at the VA's option.
- Access to electronic models of equipment being proposed is beneficial. Because electronic models will be used with CAD and visual design applications for datacenter planning, three dimensional models are preferred.

IMPLEMENTATION GUIDANCE

- It is preferred to have vendor-supplied total system replacements able to ship out to the site within 24 hours via overnight shipment means.

3.3.12 HARDWARE MANAGEMENT SOFTWARE

STANDARD

ID	***Secondary Attribute***	***Specification***
12	Type	Agent or agent-less server management tool accessible through SNMP

	Automated Notification	Hardware events including exceptions, diagnostics, and failures, shall be exposed through SNMP with documented MIBs
	Shared Repository	Allows uploading all information collected by the tool into shared repository using ODBC or JBDC in a manner that preserves data fields for the diagnostic information (e.g., date/time, error code, module, description, system identification, status, system configuration)
	Support Level	Support level of management software must match support level of server hardware

EXPLANATION OF STANDARD

The tool should allow management of the hardware through a variety of means that interface with SNMP, and must be fully documented with MIBs for all relevant data (hardware events). Integration with VA's Enterprise Exception Log Service (EELS) is a consideration and is enabled through SNMP interfacing.

Ability to access event logs through database connectivity (JDBC or ODBC) is needed for future Enterprise Management Framework tools. Service could be provided through BIOS feature, OS-level features, or other layered features, and must provide the MIB documentation for accessing the information.

EVALUATION FACTORS

- Ability to integrate with a variety of industry enterprise-level system management tools is beneficial.
- Inclusion of an enterprise-level component with shared repository and correlation analysis features is beneficial.
- Ability to provide management of agent-less systems is beneficial.

IMPLEMENTATION GUIDANCE

- Support for SNMPv3 is preferred.
- A solution design that provides secure connections outside the physical facility is preferred.
- Automated vendor notification of failure, or pre-failure activity, to initiate vendor remediation is preferred.

3.3.13 RACK INFRASTRUCTURE COMPATIBILITY

STANDARD

ID	*Secondary Attribute*	*Specification*
13	Rack Unit Measurement	Up to 4U for rack-mount servers; blade implementations will be evaluated on an individual basis.
	Rail Type	Tool-less square-hole sliding
	Cable Management	Side-reversible for non-blade solutions

EXPLANATION OF STANDARD

Rack Unit Measurement – For this class of server, occupying more than 4U (or more than two blade slots) should not be necessary; 1 slot in a half-height blade chassis is fairly equivalent to a 1U rack-mount and 1 slot in a full-height blade chassis is fairly equivalent to a 2U rack-mount. It is expected that this class of server will be a 4U rack mount or double slot full-height blade server.

Rail Type – server must have rail adapters if needed to fit the generic non-proprietary type of rack specified in this standard.

Rack infrastructure standards must match the data center equipment rack standards.

EVALUATION FACTORS

- Lower occupied rack space or blade slot (per workload) is beneficial

IMPLEMENTATION GUIDANCE

3.3.14 OPERATING SYSTEM

STANDARD

ID	*Secondary Attribute*	*Specification*
14	Type	Linux preferred based on federal requirement to utilize open source operating systems, Windows acceptable, all must VA approved operating systems and compatible with hardware.
	Version	Current target version of Linux or Windows per VA TRM
	Vendor-Installed	Depends on data center policy; default is factory-installed with default settings and most recent patch level
	Support Level / Term	Must match hardware service level; term is three years with two option years unless VA has a standing enterprise software support agreement pertinent OS.

EXPLANATION OF STANDARD

Type – Linux is preferred as the most “open” type of option, but Windows is also acceptable as an industry-wide supported option. If UNIX is required Solaris is the preferred UNIX product. Solaris is preferred due to the “open” features and high level of government security certification.

Version – Operating system build, distribution, and version must be listed in the VA TRM.

Vendor-Installed – Data center policy may require local installation of the operating system. If this is not the case, the operating system should be factory-installed and have the most recent patches applied.

Support Level / Term – The operating system support level must be consistent with the criticality of the server and match hardware service level; term is three years with two option years unless VA has a standing enterprise software support agreement pertinent OS.

EVALUATION FACTORS

- Hardware support for the latest version of OS is beneficial.
- Installation of VA-supplied image is beneficial.

IMPLEMENTATION GUIDANCE

- UNIX requires exception justification and Solaris is the preferred UNIX product.
- Preference is for the latest versions of the operating system that are certified for implementation on the VA production network and supported by the hosted applications. If earlier versions are required due to compatibility limitations, these limitations should be documented for justification.

3.3.15 PLATFORM SOFTWARE / LICENSE

STANDARD

ID	*Secondary Attribute*	*Specification*
15	Multipath	Provides automated load balancing (all paths active) and fail-over; compatible with target storage array
	Data Backup	(optional) Matches requirements of the deployment location data center and SAN backup solution
	Monitoring Tool	Matches requirements of the deployment location data center
	Layered Framework	Defined by the application owner – meet license requirements for this server
	Server Management Tools	Matches requirements of the deployment location data center
	Availability / Recovery Tools	Matches the requirements for Persistent Computing standards applied to the hosted applications/services
	Security Tools	Follows Office of Information and Technology standards including compliance with FIPS, FISMA, and VA Directives.
	Virtual Management Software	For virtual environments, an Enterprise management tool is required

EXPLANATION OF STANDARD

Multipath – Require failover capability with all paths active; Include requisite multipath software from appropriate storage vendor to provide failover and load-balancing capabilities.

Data Backup – Although this type of server primarily uses backup tools for server image backup, there may (optionally) be a need for data backup tools. A data backup tool/agent might be required by some hosted applications that have local database storage. In addition, some applications might have data stored on the SAN and is covered by the SAN backup solution that requires an agent on the server. For these servers, enterprise data backup tools are being evaluated and might lead to an enterprise standard. In the meantime, the tool selection is driven by each data center location, and must adhere to security/encryption policies. This software is installed and configured by VA and might utilize existing/enterprise licenses.

Monitoring Tool and Server Management Tool – These are the agents required by server monitoring & management tools or other monitoring tools – if agents are required. It is expected that these tools will largely be defined by the Enterprise Management Framework. The tools must meet the VA EMF standards/tools, or be able to feed information directly into the EMF tools. Until such standards are available, these tools must meet the data center standards where the tools are being placed. This software is installed and configured by VA and might utilize existing/enterprise licenses.

Availability / Recovery Tools – These include agents or other software for image backup and restoration tools, service replication tools, and fail-over tools. These will be defined in the Persistent Computing standards and by the hosted application's model for persistent computing. This software might be installed/configured through a service contract associated with the server purchase, or might be installed/configured by VA.

Security Tools – must follow Office of Information and Technology standards including compliance with FIPS, FISMA, and VA Directives, and be implemented in a way that meets Certification and Accreditation (C&A) requirements. VA Enterprise Security Solution Service (ESSS) can provide a review of proposed security tools and can be reached by contacting the Enterprise Solutions Security Service under OI&T Field Security Operations.

Virtual Management Software – Virtual infrastructure requires and enterprise level management tool which may require additional licenses for each server. Live host migration may require specific network configurations or compliance with datacenter standards. Refer to LAN Standards, Section xx and Datacenter Standards, Section xx.

EVALUATION FACTORS

- Multipath – compatibility with a wide range of storage arrays / manufacturers is a benefit.

IMPLEMENTATION GUIDANCE

3.3.16 DATACENTER MANAGEMENT

STANDARD

ID	*Secondary Attribute*	*Specification*
16	Remote Keyboard/Video/Mouse Control	Ability to remotely connect to the server's external keyboard, video, and mouse ports through a networked KVM switch.

	Remote Power Disconnect	Ability to remotely disconnect/reconnect server to power source

EXPLANATION OF STANDARD

Remote Keyboard/Video/Mouse Control – A dongle to attach the server to a networked KVM switch is required. Refer to the Datacenter Standard, Section xx

Remote Power Disconnect – Ability to remotely disconnect/reconnect power source to server through the control of individual receptacles. Refer to Datacenter Standards, Section xx.

EVALUATION FACTORS

- Ability to remotely update device firmware is beneficial.
- Accessibility via web browser is beneficial – should support Microsoft Internet Explorer.

IMPLEMENTATION GUIDANCE

3.4 VISTA BACK-END (BE) (**CLASS E**)

The VistA Back-End (BE) server will provide all Caché database data access (via SAN based storage), with data servicing to Front-End (FE) servers over redundant private network paths. The BE servers will support 1-2 Vista database Server configurations, as well as VistA taskman operations. The BE server will also host the operating system based synchronous data replication, to facilitate the High Availability requirements of the metro-based platform. The metro-based platform will consist of clustered systems to further facilitate the High Availability survivability of the overall BE Regional platform.

The provided synchronous data replication, with server clustering, for the BE servers will initially be based on the OpenVMS based operating system from HP. The OpenVMS operating system, with clustering, and in many cases shadowing, have been supporting the VistA platform since the 1990's, and is also currently deployed in the existing RDP environment for both Region 1 and Region 4. This proven operating environment will continue for the BE servers to be deployed for Regions 2 & 3 VistA consolidation.

The BE servers may also include a "quorum" node. Cluster quorum configuration options include:

- Quorum processing site
- Quorum only site (first recommendation)
- Quorum cage (second recommendation)
- Primary datacenter (second recommendation)
- Primary side, floating between datacenters

A third computer room hosting the quorum node is the preferred and recommended configuration, in order to fully support automatic High Availability failover. The secondary recommendation is to host the quorum node at one of the two metro datacenters (per region), meaning, one of the two metro datacenters will be selected as the preferred (or primary), most

survivable datacenter. As such, this datacenter will house the quorum node, meaning the VMS cluster will only survive (without manual intervention) in this datacenter, in the event of an inter-site network link failure or a datacenter failure (by definition, the only datacenter which could fail and processing continue would be the secondary datacenter).

Quorum participation can be provided by:

- Quorum node (recommended)
- Quorum disk

The Quorum node is the preferred and recommended quorum participation method. This means an additional node, of minimal configuration, will be hosted at the primary datacenter location. This node will provide the vote tie-breaking mechanism, in the event there is an inter-site network link issue.

(add text on primary recommendation sizing)

Accepting the secondary recommendation (hosting quorum at the Primary datacenter), the quorum node is then proposed to be configured as a normal BE server, but to be used for Backup processing (the processing of nightly data backups) and as an online spare in the event of longer term node outages.

3.4.1 PROCESSOR

STANDARD

ID	*Secondary Attribute*	*Specification*
1	Type	IA_64 (Itanium)
	Socket Count	Minimum required to meet Performance Benchmark standard
	Core Count	Amount required to meet Performance Benchmark standard
	Cache	Minimum required to meet Performance Benchmark standard
	Virtualization Aware	Not Required

EXPLANATION OF STANDARD

Type – The current back-end architecture is built on, and only certified on the OpenVMS operating system. OpenVMS is supported on AlphaServer and Itanium (IA_64) based systems, and Alpha based servers are no longer available.

Socket Count – OpenVMS licensing is based on core-count. Socket Count should be the minimum required to satisfy the Performance Benchmark, with the assumption that core counts will increase with future generations of supported chip types.

Core Count – minimum amount expected for performance/capacity level is dual-core, more can be better. However, any amount that meets the Performance Benchmark specification is sufficient.

Cache – Minimum required to meet Performance Benchmark standard

Virtualization Aware – Not Required

EVALUATION FACTORS

- Currency of processor – most recent speed of processor is beneficial.

IMPLEMENTATION GUIDANCE

3.4.2 PERFORMANCE

STANDARD

ID	*Secondary Attribute*	*Specification*
2	Performance Benchmark	SPEC CPU2006 Integer Rate Base value >= 90

EXPLANATION OF STANDARD

This SPEC_Int rate is represented by a large number of server platforms. Performance requirement is largely dependent on the application's workload and compute-resource consumption profile. For example, the Cleveland VAMC is considered a very large VistA environment, and has been piloting the Itanium based server Front End. The daily average M Command utilization for Cleveland is reported to be 800,000 to 900,000. These levels consume the older ES80 model 8 (8 CPUs), and utilize about 1/3 to ½ of an Itanium based server (for the daily average). Other data suggests Itanium based servers are more than twice as powerful as Alpha based servers, with similar core or CPU counts (as appropriate for the technology). Considering the peak period requirements, and allowing for performance spikes, as well as the failover and switchover requirements will generally double the capacity requirements of the BE severs, implementing servers twice as capable as the currently deployed alpha platform is appropriate. The SPEC2000 value for deployed ES80 systems is 62 or 73 (depending on chip speed). The SPECint_rate2006 (published Nov 2006) value for rx6600/BL870c (Itanium based) servers is 91.1. It is assumed the Spec2006 performance mtric is more difficult, and returns lowers values, as compared with SPEC2000 ratings. For this exercise, it is assumed SPEC2006 values would be about ½ what they might have been if measured under SPEC2000 (i.e. the rx6600 SPEC2006 value of 91.1 might have had a SPEC2000 value near 180).

EVALUATION FACTORS

- Higher benchmark rating value is beneficial.

IMPLEMENTATION GUIDANCE

3.4.3 MEMORY

STANDARD

ID	*Secondary Attribute*	*Specification*
3	Type	Compatible with CPU with full error correction capability including ECC
	Total Memory	Minimum 24GB usable
	Speed	Minimum - match operational CPU and bus speed required to meet the Performance Benchmark requirement

EXPLANATION OF STANDARD

Type – Memory has to be highly reliable – ECC is the expected feature to allow higher levels of reliability. RDIMM is more expandable/scalable and provides additional reliability features (e.g., more error-checking features) while consuming more power (about one watt per chip) and costing more. In general, RDIMM is recommended with high memory requirements and higher availability requirements.

Total Memory – Itanium memory usage expected to be about 20% greater than existing alpha-based system. InterSystems' recommendation also calls for 3-5 MB per process, in terms of global buffer sizing.

Using 16 GB for a base server, with 8 cores, and adding 2.5 GB (based on 500 Taskman jobs with 5MB buffer space per process) for global buffer cache, with any server possibly supporting up to 3 configurations (in a failover or switchover situation), this sizing accounts for 16 GB for typical server processing (Db serving, Taskman processes, general overhead) plus 7.5 GB for Caché global buffers.

Speed – Memory speed is expected to be matched by the vendor, to match processor capabilities. As Itanium servers are expected to be only available from HP, vendor dependence is warranted.

EVALUATION FACTORS

- Additional memory is beneficial
- Higher density chips while avoiding a decrease in memory bus speed is preferred.
- Speed of memory takes precedence over density as long as capacity is met

IMPLEMENTATION GUIDANCE

- RDIMM is recommended for high memory requirements and higher availability requirements.

3.4.4 INTERNAL STORAGE

STANDARD

ID	*Secondary Attribute*	*Specification*

4	Type	SAS
	Quantity of Drives	Minimum one (DOSD – i.e. to host Dump Off System Disk)
	RAID Configuration	Not Required
	Drive Speed	Minimum 10K RPM
	Drive Capacity	Minimum 36GB usable storage capacity
	Disk Subsystem Performance	The disk must meet or exceed the performance levels specified below for tests "Transaction Rate" and "Throughput Rate" Validation will be incorporated in user acceptance phase.
	Controller Cache	Not Required
	Quantity of Controllers	Minimum of one

EXPLANATION OF STANDARD

Type – SAS is required for the following reasons: 1) SAS uses the SCSI command set which provides better error reporting and recovery techniques than does SATA, 2) SATA throughput can be degraded on SAS adapters that use the Serial Tunneling Protocol (STP) to communicate with SATA targets, and 3) because they are engineered for continuous duty cycles and longer MTBFs, SAS drives are an accepted industry standard for enterprise server solutions.

Quantity – The Backend servers will SAN boot, and otherwise access all required data from the SAN. However, a local disk for possible booting, or more likely hosting the creation of system crash memory dumps will be useful. A local disk may be required for this function.

RAID Configuration – An internal disk could serve as a host to operating system dumps generated thru a system crash. The longer term retention of these files is also considered. Assuming dump files of 10 GB in size (or less), the desire is to retain at least 3 copies of recent crash dumps. An internal system disk is not expected to be required.

Drive Speed – Minimum RPM for this class of server & storage is 10K..

Drive Capacity – 36GB minimum usable storage capacity is sufficient for the dump file, and optional older dump file disk storage needs.

Disk Subsystem Performance – The disk must meet or exceed the performance levels specified below for tests "Transaction Rate" and "Throughput Rate" as described in APPENDIX G.

- Transaction Rate = 500+ IOPS
- Throughput Rate = 50+ MBPS

Controller Cache – Cache is not required because disk subsystem will only be used to store dump files.

Quantity of Controllers – Only one SAS controller is required.

EVALUATION FACTORS

- Higher storage capacity is beneficial
- Higher RPM is beneficial for drive platters
- Additional internal drives is beneficial

IMPLEMENTATION GUIDANCE

3.4.5 EXTERNAL STORAGE INTERFACE

STANDARD

ID	*Secondary Attribute*	*Specification*
5	Type	Fiber-channel with duplex LC connectors
	Port Count	Minimum of four ports with full port redundancy provided.
	Bandwidth	Minimum 4Gbps; must meet or exceed SAN fabric requirements
	Compatibility	Support storage arrays at deployment location; supported by operating system and multipath drivers specified for this server
	SAN Boot Capability	BIOS support for SAN boot
	Blade Chassis Component (if applicable)	Pass thru connections, for increased bandwidth capabilities
	Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the SAN. Minimum redundant ports

EXPLANATION OF STANDARD

- Type – fiber channel and Ethernet are the vast majority of investment for connections used in storage-dedicated networks in VA data centers, and exceptions should be justified including impact to existing infrastructure investment. Typical Ethernet connector types are RJ-45 for copper and LC for optical. When Ethernet is used, TCP Offload Engine (TOE) is preferred for all servers. Drivers are more readily available for Windows than Linux servers.
- Port Count – minimum required for port-level redundancy / load-balancing. Provides the flexibility to improve the throughput and performance without further hardware changes. Can initially configure with only 2 active.

Potential benefit for separating the ports used for booting and data. Allows increased workload of multiple database servers in shared server environment.

- Bandwidth – 4Gbps is standard for current investment of optical fabrics in the storage systems and infrastructure in VA data centers and matches the minimum current industry standard; similarly 1Gbps is standard for current investments and minimum current industry practices when using copper-based storage fabrics.
- Compatibility – Support storage subsystems at deployment locations. Check adapter compatibility for storage at deployment location. Don't want to provide a list of brand names.
- Blade Chassis Components – Considering the bandwidth requirements of network connections (Ethernet and SAN), pass-thru connections preferred at this time.
- Hardware Redundancy – All storage communication connections should provide full redundancy, from the option card, port, and switch connectivity.

EVALUATION FACTORS

- For a server with high criticality of SAN storage connectivity, redundant adapter cards are recommended.
- Interface for iSCSI can be termed an HBA interface or a NIC interface. Where higher performance is a concern, the HBA interface is recommended.

IMPLEMENTATION GUIDANCE

- For a server with high criticality of SAN storage connectivity, redundant adapter cards are recommended.
- Interface for iSCSI can be termed an HBA interface or a NIC interface. Where higher performance is a concern, the HBA interface is recommended.
- Bandwidth spread across multiple ports on independent cards may be required for performance reasons
- Backend Server for VistA require daily disk to disk backups of both synchronous copies of production databases as well as for both copies of the CSJ provided databases (between regions). Further, at least one copy of each database (production and CSJ) must be archived for later restoration if needed. The Production and CSJ database backups strategy should be documented in the SAN/Storage/Back-up standards document.

3.4.6 SERVER COMMUNICATIONS INTERFACE

STANDARD

ID	*Secondary Attribute*	*Specification*
6	Type	Ethernet matching connector type for network in data center
	Port Count	Minimum of 6 ports
	Bandwidth	Minimum 1Gbps
	Hardware Redundancy	Supports connecting redundant ports configured to separate switches on the LAN; supports physically independent redundant ports including embedded ports for each interface function (e.g., public network function and cluster interconnect function)

	Blade Chassis Component (if applicable)	Pass Thru connections

EXPLANATION OF STANDARD

Type – Ethernet is the vast majority of investment for connections used in data center networks in VA, and exceptions should be justified including impact to existing infrastructure investment. Typical Ethernet connector types are RJ-45 for copper and LC for optical with a future trend towards copper only.

Port Count – when used in a virtual environment, a base set of 4 ports are used for virtual guests, 1 is for the console, and 1 for the live migration. Other specific interfaces may be required depending on the profile of the workload/application. When not used in a virtual environment, cluster interconnects may need 2 or more ports, depending on the amount of cluster traffic (e.g., Oracle clusters send data over the cluster interconnect, simple Microsoft clustering only send "heartbeat" traffic).

Bandwidth – expected standard interface on network switches is 1Gbps. There is a potential for 10Gbps Ethernet.

Hardware Redundancy – embedded ports count as one "card"; all interface uses are spread across two "cards" (including at least one add-on NIC). The goal is physical separation of the chipset/ASIC used on the redundant port, for each function. Support for "simple failover" (with a single IP address) and independent usage (with multiple IP addresses) is required.

Blade Chassis Component – Due to the small form factor of the blade chassis, pass thru connections are required.

EVALUATION FACTORS

- Additional cards can be beneficial
- Support for active/active load balancing is beneficial

IMPLEMENTATION GUIDANCE

- Backend Server for VistA requires ECP (Caché Enterprise cache Protocol) communication between the Backend DB servers, and Frontend Application Servers. This ECP communication must support the spanning of metro-datacenters, without changing the underlying IP address, as defiend by the Application Server connection to the DB server. The Production ECP communication strategy should be documented in the LAN standards document (see GSS/ACE discussion).

- Backend Server for VistA requires SCS (VMScluster System Communications Services) communication between each cluster member in both (all) metro-datacenters. The SCS protocol cannot be routed, and thus must be bridged. The Production SCS communications strategy should be documented in the LAN standards document (see SCS ACE or layer 2 tunneling discussion).

3.4.7 REMOVABLE MEDIA DEVICE

STANDARD

ID	*Secondary Attribute*	*Specification*
7	Optical Drive	CD/DVD drive – minimum 24x speed
	USB	Not Required

EXPLANATION OF STANDARD

The functional requirement supported by this standard is to provide a media transfer mechanism to upload patch and software/firmware installation, upload diagnostic tools, and (with read/write capability) to download diagnostic information. Manufacturers are going toward a trend of using USB flash drives for sending patch/upgrade/firmware distributions. Therefore, the requirement is to provide both CD/DVD and USB flash drive capabilities. Note that security restrictions or concerns need to be addressed before using removable media.

The type of drive specified is the typical expected minimum specification and type.

Alternative Removal Media Device access must be quoted and provided in cases where an internal device of this type is not configured or available.

EVALUATION FACTORS

- Having a DVD drive is beneficial.
- Faster speeds and/or Read-write capability is beneficial (additional security checks may be required for read/write capability).

IMPLEMENTATION GUIDANCE

- Preferred to provide removable media device for patch and software installation options; but small form factor (e.g., blade) servers might use other options. Media access type must be provided for, as configured by vendor.

3.4.8 POWER

STANDARD

ID	*Secondary Attribute*	*Specification*
8	Redundancy	N+1 redundancy on power supply
	Voltage	208V input voltage on power supply
	Cord Type	IEC

	Power Efficiency	Until such time that there are Energy Star ratings available to most server classes, power efficiency should be consistent with industry trends.
	Power Consumption Documentation	Vendor-supplied data for power consumption of specified configuration at 100% load, “typical” active usage, and at idle.

EXPLANATION OF STANDARD

Redundancy – server must continue to operate if and when a power supply fails in order to enhance availability.

Voltage and cord type – driven by data center standards, 208V can more efficiently use power from utility

Power Efficiency – Certifications for power efficiency include: 80-plus (Dell & HP certify through them, but IBM does not), and Energy Star (still forming the standard for servers, but blade servers will be separate standard). The EPA establishes Energy Star standards – tier 2 appears to be the four-socket type of server while tier 1 appears to be the dual-socket server types. Some server manufacturers implement power efficiency management features including step-down of power consumption during low resource-consumption states.

Power Consumption Documentation – Preferred documentation is Energy Star data sheet and/or SPECpower_ssj2008 published results.

EVALUATION FACTORS

- Power management/economy and reporting/analysis features are beneficial.
- Lower power consumption rating at 100% utilization per workload capacity (in terms of SPECint_rate value) is beneficial.

IMPLEMENTATION GUIDANCE

- Cord type of IEC-C13 is preferred for most classes of rack-mount servers

3.4.9 COOLING FAN

STANDARD

ID	***Secondary Attribute***	***Specification***
9	Redundancy	Minimum N+1 redundancy

EXPLANATION OF STANDARD

Redundancy – server must continue to operate if and when a fan fails in order to enhance availability.

EVALUATION FACTORS

- Additional fans or separation of fans into multiple cooling zones is beneficial.
- Lower total decibel rating on the fans is beneficial.

IMPLEMENTATION GUIDANCE

3.4.10 OUT OF BAND MANAGEMENT

STANDARD

ID	*Secondary Attribute*	*Specification*
10	Type	Remote command-line and console-level access utilizing dedicated network interface
	Security	Secure IP-based remote management that complies with all VA security requirements
	Remote Power Control	Ability to remotely power on/off/reset server

EXPLANATION OF STANDARD

Required dedicated network interface is required outside of the other network interfaces specified in section 3.2.7.

Based on data center environment, operational needs, and minimum requirements that are typically satisfied through a wide variety of servers. Includes access for diagnostic capability.

Remote Power Control – Ability to power on/off/reset through server based management.

Console logging – OpenVMS generates OPCOM (Operator Communications) messages related to operating system related actions and possible error conditions. This information can be invaluable after an event, as part of troubleshooting. It is required that a solution be offered to log all OPCOM messages from the console, for historical purposes.

EVALUATION FACTORS

- Ability to remotely mount a drive or ISO is beneficial. This can be used for software, firmware and OS updates using remote drive mounting, and it minimizes the need for local optical drive.
- Active Directory integration for authentication is beneficial.
- Role-based access control is beneficial.
- Accounting capability is beneficial.
- Strong password compliance is beneficial.
- Accessibility via web browser is beneficial – should support Microsoft Internet Explorer.
- Support for Secure Shell (SSH) connections is beneficial – should be FIPS 140-2 compliant.
- Ability to remotely identify server with visual indicator (e.g., indicator light) is beneficial.

IMPLEMENTATION GUIDANCE

3.4.11 SERVER HARDWARE VENDOR SUPPORT

STANDARD

ID	*Secondary Attribute*	*Specification*
11	Type	Match service level requirement for hosted applications; mission-critical level support shall include 24x7 coverage 365 days/year, less than 1 hour initial engineer-level response time, 4-hour on-site response time for emergency dispatches, and 8-hour time-to-repair
	Term Length	Minimum of three years with two additional option years
	Keep Your Hard Drive	All storage drives are kept by VA in the event of warranty/service replacement

EXPLANATION OF STANDARD

Type – the level of support matches the criticality level of the hosted services. Mission-critical level is defined in this standard, and has been shown to be accommodated by major server vendors. Other levels will have to be defined to match the business expectations of that criticality level. A non-critical example would be 8 business hours/weekday phone coverage with next-day response and repair. Access to original server manufacturer support with the option to escalate to senior technicians at VA discretion is required.

Term – some components might only allow three years of support on an initial contract, but option years must be available to extend that to the expected lifetime of the system which is nominally five years.

Keep Your Hard Drive – This is a security requirement that is now a de-facto standard in VA.

EVALUATION FACTORS

- Lower on-site response time is beneficial.
- Lower initial “engineer level” contact response time is beneficial.
- Lower time-to-repair is beneficial.
- Vendor-supplied components stocked as spares near, or at, facility is beneficial.
- Penalty clauses for vendor not meeting support contract provides benefit.
- Allowance for optional parts-delivery/customer-installation of hot swappable components, memory, and drop-in cards beneficial at the VA’s option.
- Access to electronic models of equipment being proposed is beneficial. Because electronic models will be used with CAD and visual design applications for datacenter planning, three dimensional models are preferred.

IMPLEMENTATION GUIDANCE

- It is preferred to have vendor-supplied total system replacements able to ship out to the site within 24 hours via overnight shipment means.

3.4.12 HARDWARE MANAGEMENT SOFTWARE

STANDARD

ID	*Secondary Attribute*	*Specification*
12	Type	Agent or agent-less server management tool accessible through SNMP
	Automated Notification	Hardware events including exceptions, diagnostics, and failures, shall be exposed through SNMP with documented MIBs
	Shared Repository	Allows uploading all information collected by the tool into shared repository using ODBC or JBDC in a manner that preserves data fields for the diagnostic information (e.g., date/time, error code, module, description, system identification, status, system configuration)
	Support Level	Support level of management software must match support level of server hardware

EXPLANATION OF STANDARD

The tool should allow management of the hardware through a variety of means that interface with SNMP, and must be fully documented with MIBs for all relevant data (hardware events). Integration with VA's Enterprise Exception Log Service (EELS) is a consideration and is enabled through SNMP interfacing.

Ability to access event logs through database connectivity (JDBC or ODBC) is needed for future Enterprise Management Framework tools. Service could be provided through BIOS feature, OS-level features, or other layered features, and must provide the MIB documentation for accessing the information.

EVALUATION FACTORS

- Ability to integrate with a variety of industry enterprise-level system management tools is beneficial.
- Inclusion of an enterprise-level component with shared repository and correlation analysis features is beneficial.
- Ability to provide management of agent-less systems is beneficial.

IMPLEMENTATION GUIDANCE

- Support for SNMPv3 is preferred.
- A solution design that provides secure connections outside the physical facility is preferred.
- Automated vendor notification of failure, or pre-failure activity, to initiate vendor remediation is preferred.

3.4.13 RACK INFRASTRUCTURE COMPATIBILITY

STANDARD

ID	*Secondary Attribute*	*Specification*
13	Rack Unit Measurement	7U or 10U for rack-mount servers; 4-slot blade for blade servers
	Rail Type	Tool-less square-hole sliding
	Cable Management	Side-reversible for non-blade solutions

EXPLANATION OF STANDARD

- Rack Unit Measurement – For this class of server, occupying more than 10U (or four blade slots) should not be necessary
- Rack infrastructure standards must match the data center equipment rack standards.

EVALUATION FACTORS

- Lower occupied rack space or blade slot (per workload) is beneficial

IMPLEMENTATION GUIDANCE

3.4.14 OPERATING SYSTEM

STANDARD

ID	*Secondary Attribute*	*Specification*
14	Type	OpenVMS, Un-Limited User Licensing, VMS clustering, Host Based Volume Shadowing, TCP/IP Services
	Version	OpenVMS version 8.3 (minimum version support)
	Vendor-Installed	Not Required, but preferred
	Support Level / Term	Must match hardware service level; term is three years with two option years unless VA has a standing enterprise software support agreement pertinent OS.

EXPLANATION OF STANDARD

Type – OpenVMS based servers are the current and preferred operating system for Back End server to support the replication and availability requirements of the RDP consolidated environments.

Unlimited User licensing will offer the flexibility of supporting additional users as needed for automatic failover operations (when nodes or data centers fail) as well as switchover operations, when a regions entire metro-area cluster fails.

VMS clustering will facilitate and allow automatic node, configuration, and VistA failover, in the event of node failures or crashes. Clustering will also help facilitate failover if and when single data-centers fail or otherwise become unavailable for operations.

Host Based Volume Shadowing (HBVS) is the supported mechanism of OpenVMS to synchronously replicate data between metro-data centers assuring data survivability against metro data center failure or loss.

TCP/IP is the currently deployed TCP solution for OpenVMS, within the VA.Version – Operating system build, distribution, and version must be listed in the VA TRM.

Vendor-Installed – Not Required, but preferred

Support Level / Term – The operating system support level must be consistent with the criticality of the server and match hardware service level; term is three years with two option years unless VA has a standing enterprise software support agreement pertinent OS.

EVALUATION FACTORS

- Hardware support for the latest version of OS is beneficial.

IMPLEMENTATION GUIDANCE

3.4.15 PLATFORM SOFTWARE / LICENSE

STANDARD

ID	*Secondary Attribute*	*Specification*
15	Multipath	Provides automated load balancing (all paths active) and fail-over; compatible with target storage array
	Data Backup	(optional) Matches requirements of the deployment location data center and SAN backup solution
	Monitoring Tool	Matches requirements of the deployment location data center
	Layered Framework	Support for InterSystems' Caché (minimum version 5.2.3)
	Server Management Tools	Matches requirements of the deployment location data center
	Availability / Recovery Tools	Matches the requirements for Persistent Computing standards applied to the hosted applications/services

	Security Tools	Follows Office of Information and Technology standards including compliance with FIPS, FISMA, and VA Directives.
	Virtual Management Software	Not Required

EXPLANATION OF STANDARD

Multipath – Require failover capability with all paths active; Include requisite multipath software from appropriate storage vendor to provide failover and load-balancing capabilities.

Data Backup – Although this type of server primarily uses backup tools for server image backup, there may (optionally) be a need for data backup tools. A data backup tool/agent might be required by some hosted applications that have local database storage. In addition, some applications might have data stored on the SAN and is covered by the SAN backup solution that requires an agent on the server. For these servers, enterprise data backup tools are being evaluated and might lead to an enterprise standard. In the meantime, the tool selection is driven by each data center location, and must adhere to security/encryption policies. This software is installed and configured by VA and might utilize existing/enterprise licenses.

Monitoring Tool and Server Management Tool – These are the agents required by server monitoring & management tools or other monitoring tools – if agents are required. It is expected that these tools will largely be defined by the Enterprise Management Framework. The tools must meet the VA EMF standards/tools, or be able to feed information directly into the EMF tools. Until such standards are available, these tools must meet the data center standards where the tools are being placed. This software is installed and configured by VA and might utilize existing/enterprise licenses.

Layered Framework – The VistA application suite is built on InterSystems' Caché, with the currently deployed version of 5.2.3.

Availability / Recovery Tools – These include agents or other software for image backup and restoration tools, service replication tools, and fail-over tools. These will be defined in the Persistent Computing standards and by the hosted application's model for persistent computing. This software might be installed/configured through a service contract associated with the server purchase, or might be installed/configured by VA.

Security Tools – must follow Office of Information and Technology standards including compliance with FIPS, FISMA, and VA Directives, and be implemented in a way that meets Certification and Accreditation (C&A) requirements. VA Enterprise Security Solution Service (ESSS) can provide a review of proposed security tools and can be reached by contacting the Enterprise Solutions Security Service under OI&T Field Security Operations.

Virtual Management Software – Virtual infrastructure requires and enterprise level management tool which may require additional licenses for each server. Live host migration may require specific network configurations or compliance with datacenter standards. Refer to LAN Standards, Section xx and Datacenter Standards, Section xx.

EVALUATION FACTORS

- Multipath – compatibility with a wide range of storage arrays / manufacturers is a benefit.

IMPLEMENTATION GUIDANCE

3.4.16 DATACENTER MANAGEMENT

STANDARD

ID	*Secondary Attribute*	*Specification*
16	Remote Keyboard/Video/Mouse Control	Ability to remotely connect to the server's external keyboard, video, and mouse ports through a networked KVM switch.
	Remote Power Disconnect	Ability to remotely disconnect/reconnect server to power source

EXPLANATION OF STANDARD

Remote Keyboard/Video/Mouse Control – A dongle to attach the server to a networked KVM switch is required. Refer to the Datacenter Standard, Section xx

Remote Power Disconnect – Ability to remotely disconnect/reconnect power source to server through the control of individual receptacles. Refer to Datacenter Standards, Section xx.

EVALUATION FACTORS

- Ability to remotely update device firmware is beneficial.
- Accessibility via web browser is beneficial – should support Microsoft Internet Explorer.

IMPLEMENTATION GUIDANCE

4 TAXONOMY OF STANDARDS

[Build taxonomy of this standard, defining the primary and secondary attributes common to all standards sets that will be contained in this standard. Add additional rows and cells as necessary. This attribute set will be used to populate the tables in Sections 2 and 3. Replace this paragraph with a paragraph introducing your taxonomy.]

ID	*Primary Attribute*	*Secondary Attribute*
1	Processor	Type

ID	*Primary Attribute*	*Secondary Attribute*
		Socket Count
		Core Count
		Cache
		Virtualization Aware
2	Performance	Performance Benchmark
3	Memory	Type
		Total Memory
		Speed
4	Internal Storage	Type
		Quantity of Drives
		RAID Configuration
		Drive Speed
		Drive Capacity
		Disk Subsystem Performance
		Controller Cache
		Quantity of Controllers
5	External Storage Interface	Type
		Port Count
		Bandwidth
		Compatibility
		SAN Boot Capability
		Blade Chassis Component (if applicable)
		Hardware Redundancy
6	Server Communications Interface	Type
		Port Count
		Bandwidth
		Hardware Redundancy
		Blade Chassis Component (if applicable)
7	Removable Media Device	Optical Drive

ID	*Primary Attribute*	*Secondary Attribute*
		USB
8	Power	Redundancy
		Voltage
		Cord Type
		Power Efficiency
		Power Consumption Documentation
9	Cooling Fan	Redundancy
10	Out of Band Management	Type
		Security
		Remote Power Control
11	Server Hardware Vendor Support	Type
		Term Length
		Keep Your Hard Drive
12	Hardware Management Software	Type
		Automated Notification
		Shared Repository
		Support Level
13	Rack Infrastructure Compatibility	Rack Unit Measurement
		Rail Type
		Cable Management
14	Operating System	Type
		Version
		Vendor-Installed
		Support Level / Term
15	Platform Software / License	Multipath
		Data Backup
		Monitoring Tool

ID	*Primary Attribute*	*Secondary Attribute*
		Layered Framework
		Server Management Tools
		Availability / Recovery Tools
		Security Tools
		Virtual Management Software
16	Datacenter Management	Remote Keyboard/Video/Mouse Control
		Remote Power Disconnect

APPENDIX A – DEFINITIONS

- **Mission-Critical** - One description of mission-critical criticality level for system platforms is provided below for reference:

  Mission Critical Information
  - Business Value: Essential to fundamental business operations – outage seriously impairs functioning of business
  - Outage Impact: Interruption of service may result in severe financial, regulatory, safety, patient health, or other business issues
  - Significant Outage: More than five minutes of downtime is considered significant at any time

- **Cache** - Small memories on or close to the CPU that can speed up processing by reducing the number of times a CPU has to access much slower main memory. All CPUs have level 1 (L1) cache and may also have level 2 (L2) and level 3 (L3) cache.

- **SPEC CPU2006 Integer Rate Base** - CPU2006 is a set of benchmarks designed to test the CPU performance of a server computer system. This benchmark is usually run on only a single CPU (even if the system has many CPUs) and only on a single core if the CPU has multiple cores. The "SPEC CPU2006 Integer Rate Base" is a more complete system-level benchmark that allows all CPUs to be used.

- **N+1** - Need plus one, a redundancy concept where capacity is configured to include planned capacity plus one additional device to enable continued operations with the failure of one system in the configuration. This presumes immediate detection and remediation of the failed unit.

- ***n*U (or Rack Unit)**- A rack unit or U is a unit of measure used to describe the height of equipment intended for mounting in a 19-inch rack. The size of a piece of rack mounted equipment is frequently described as a number in "U". For example, one rack unit is often referred to as "1U", 2 rack units as "2U" and so on.

- **SAS (Serial Attached SCSI)** - A data-transfer technology that moves data to and from computer storage devices such as hard drives and tape drives.

- **ECC (Error-Correcting Code)** - Due to electrical or magnetic interference inside a computer system, single bits of memory can spontaneously flip to the opposite state. To mitigate this, ECC enabled memory detect single-bit errors and correct them.

- **CPU (Central Processing Unit)** - The Central Processing Unit (CPU) or processor is the portion of a computer system that carries out the instructions of a computer program, and is the primary element carrying out the computer's functions.

- **x86_64** - In CPU design, x86-64 is an extension of the x86 instruction set. It allows far larger virtual and physical address spaces than x86_32 designs, doubles the width of the integer registers from 32 to 64 bits, increases the number of integer registers, and provides other enhancements. It is also known as x64.

APPENDIX B – REFERENCES

SPEC http://www.spec.org/

Serial ATA International http://www.serialata.org/index.asp

Oracle Weblogic http://www.oracle.com/appserver/weblogic/weblogic-suite.html

VA LAN Standards [place URL link here]

VA WAN Standards [place URL link here]

VA Datacenter Standards [place URL link here]

VA Storage Standards [place URL link here]

APPENDIX C – ACRONYMS

Refer to the VA Acronym Lookup Web page for a list of VA specific acronyms.

NIC	Network Interface Controller
HBA	Host Bus Adapter
DDR2	Double Data Rate 2 (a Memory interface specification)
AMD	Advanced Micro Devices, Inc.
SPEC	Standard Performance Evaluation Corporation
ECC	Error-Correcting Code
MB	Megabytes
Mbps	Megabits per Second
GB	Gigabytes
Gbps	Gigabits per Second
RPM	Rotations per Minute
IEC	International Electrotechnical Commission
GHz	Gigahertz
SAS	Serial Attached SCSI
PCI-E	PCI-Express (or Peripheral Component Interconnect Express)
RDIMM	Registered Dual In-line Memory Module
ASIC	Application-Specific Integrated Circuit
LC	Lucent Connector (or Local Connector)
BIOS	Basic Input/Output System
IOPS	Input/Output Operations Per Second

APPENDIX D – CONTRIBUTORS

Michael Julian	Data center Design
Michael Bell	Information Modeling
David Deaderick	Linux/Oracle/WebLogic Server Hardware
John Beaufait	
David Carter	
Nelda Cook	
John Mehl	
Robert Moser	
Franco Susi	
Kurt Virnig	
James Yaple	
David Wong	
Mark Zimmerman	
Denise McLain	
John Dellar	
Jeff Lawlor	

APPENDIX F – SERVER CLASSES

The following list of common server functions associates function and workload with particular server classes. This list is not intended to be all inclusive but simply a starting point or guide on how to map server function and workload with a server class.

The following table provides a mapping of server functions to server classes:

Server Function	Standard Class for LIGHT Workload	Standard Class for TYPICAL Workload	Standard Class for HEAVY Workload
AudioCare	CLASS C[1]	CLASS C[1]	CLASS C[1]
Domain Controller	CLASS C[7]	CLASS B	CLASS B
Data Backup	CLASS B	CLASS B	CLASS B
File server	CLASS C	CLASS B	CLASS B
Print server	CLASS C	CLASS C	CLASS C
Web or FTP server	CLASS C	CLASS C	CLASS C
Database	CLASS B	CLASS B	CLASS A
Exchange	N/A	CLASS A[4]	CLASS A[4]
Management Application (e.g. ePO, Ciscoworks, Orion, etc)	CLASS B[5]	CLASS A[5]	CLASS A[5]
Management Application w/o database (SAN management app – e.g. HP Commandview)	CLASS C	CLASS B	CLASS B
Application server (e.g. Weblogic JVM)	CLASS C	CLASS B	CLASS B
SCCM Application	N/A	CLASS B	CLASS A
SCCM Database	N/A	CLASS B	CLASS A
SCCM Delivery	N/A	CLASS C[8]	CLASS B
SharePoint MOSS	N/A	CLASS B	CLASS A
Terminal Server	CLASS C	CLASS B	CLASS B
Virtual Host	CLASS B[6]	CLASS A	CLASS A
VistA Front-end Application	CLASS C[3]	CLASS C[3]	CLASS C[3]
VistA Imaging Background Processor	CLASS B[2]	CLASS B[2]	CLASS B[2]
VistA Imaging Gateway	CLASS C	CLASS C	CLASS C
WSUS (update servers)	CLASS C	CLASS B	CLASS B
VistA Back-end server	N/A	CLASS E	CLASS E

1. Class C with potential need for form factor exemption
2. Class B due to clustering and SAN activity
3. Quantity dependent (used in a scale-out model)
4. Class A with application specific modifications
5. Class A or B; workload will vary depending upon whether it is dedicated or shared
6. Used to host one or two virtual guests for high availability only
7. Read-Only domain controller
8. May require external storage

APPENDIX G – DISK SUBSYSTEM PERFORMANCE TEST

The tests described in this section use IOMeter version 2006.07.27 to evaluate the performance of a server's disk subsystem. There are two tests described; the "Transfer Rate" test which is measured in I/O operations per second (IOPS) and the "Throughput Rate" test which is measured in megabytes per second (MBPS).

Before the tests can be run, IOMeter must be setup with a baseline configuration. This can be achieved by copying the contents of TABLE 1 into a file, saving the file with a .icf extension, and opening it with IOMeter. The final step in configuring IOMeter is setting the queue depth to 4. On the first screen when IOMeter is started, there is a parameter named "*# of Outstanding I/O's*". Change this parameter from 1 to 4.

Once IOMeter has been configured, perform the following:

1. Select "*Worker 1*" in the left pane.
2. Select the "*Disk Targets*" tab then select the disk to be tested.
3. Select the "*Access Specifications*" tab then select the "*4K 67% Read – Random*" test in the left pane and the "*Add*" button.
4. Select the "*Start Tests*" button (the green flag on the top toolbar). When asked for a file name, use the default value of "*results*".
5. Let the test run until it stops (5 minutes)

The results are located in the "results.csv" file create before the test. Opening the file in Excel, the test results can be found:

TEST 1: TRANSACTION RATE – Column G under "IOps"

TEST 2: THROUGHPUT RATE – Column J under "MBps"

TABLE 1: IOMETER CONFIGURATION FILE

```
Version 2006.07.27
'TEST SETUP ====================================================================
'Test Description

'Run Time
'               hours       minutes     seconds
                0           5           0
'Ramp Up Time (s)
                0
```

```
'Default Disk Workers to Spawn
                NUMBER_OF_CPUS
'Default Network Workers to Spawn
                0
'Record Results
                ALL
'Worker Cycling
'               start      step       step type
                1          1          LINEAR
'Disk Cycling
'               start      step       step type
                1          1          LINEAR
'Queue Depth Cycling
'               start      end        step       step type
                1          32         2          EXPONENTIAL
'Test Type
                NORMAL
'END test setup
'RESULTS DISPLAY ===============================================================
'Update Frequency,Update Type
                0,WHOLE_TEST
'Bar chart 1 statistic
                Total I/Os per Second
'Bar chart 2 statistic
                Total MBs per Second
'Bar chart 3 statistic
                Average I/O Response Time (ms)
'Bar chart 4 statistic
                Maximum I/O Response Time (ms)
'Bar chart 5 statistic
                % CPU Utilization (total)
'Bar chart 6 statistic
                Total Error Count
'END results display
'ACCESS SPECIFICATIONS =========================================================
'Access specification name,default assignment
                Default,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                4096,100,67,100,0,1,0,0
'Access specification name,default assignment
                4K 67% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                4096,100,67,100,0,4,0,0
'Access specification name,default assignment
                4K 100% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                4096,100,100,100,0,4,0,0
'Access specification name,default assignment
                4K 0% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                4096,100,0,100,0,4,0,0
'Access specification name,default assignment
                4K 100% Read - Sequential,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                4096,100,100,0,0,4,0,0
'Access specification name,default assignment
                4K 0% Read - Sequential,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                4096,100,0,0,0,4,0,0
'Access specification name,default assignment
                32K 67% Read - Random,NONE
```

```
'size,% of size,% reads,% random,delay,burst,align,reply
               32768,100,67,100,0,4,0,0
'Access specification name,default assignment
               32K 100% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               32768,100,100,100,0,4,0,0
'Access specification name,default assignment
               32K 0% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               32768,100,0,100,0,4,0,0
'Access specification name,default assignment
               32K 100% Read - Sequential,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               32768,100,100,0,0,4,0,0
'Access specification name,default assignment
               32K 0% Read - Sequential,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               32768,100,0,0,0,4,0,0
'Access specification name,default assignment
               64K 67% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               65536,100,67,100,0,4,0,0
'Access specification name,default assignment
               64K 100% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               65536,100,100,100,0,4,0,0
'Access specification name,default assignment
               64K 0% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               65536,100,0,100,0,4,0,0
'Access specification name,default assignment
               64K 100% Read - Sequential,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               65536,100,100,0,0,4,0,0
'Access specification name,default assignment
               64K 0% Read - Sequential,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               65536,100,0,0,0,4,0,0
'Access specification name,default assignment
               128K 67% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               131072,100,67,100,0,4,0,0
'Access specification name,default assignment
               128K 100% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               131072,100,100,100,0,4,0,0
'Access specification name,default assignment
               128K 0% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               131072,100,0,100,0,4,0,0
'Access specification name,default assignment
               128K 100% Read - Sequential,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               131072,100,100,0,0,4,0,0
'Access specification name,default assignment
               128K 0% Read - Sequential,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
               131072,100,0,0,0,4,0,0
'Access specification name,default assignment
               256K 67% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
```

```
                262144,100,67,100,0,4,0,0
'Access specification name,default assignment
                256K 100% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                262144,100,100,100,0,4,0,0
'Access specification name,default assignment
                256K 0% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                262144,100,0,100,0,4,0,0
'Access specification name,default assignment
                256K 100% Read - Sequential,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                262144,100,100,0,0,4,0,0
'Access specification name,default assignment
                256K 0% Read - Sequential,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                262144,100,0,0,0,4,0,0
'Access specification name,default assignment
                512K 67% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                524288,100,67,100,0,4,0,0
'Access specification name,default assignment
                512K 100% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                524288,100,100,100,0,4,0,0
'Access specification name,default assignment
                512K 0% Read - Random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                524288,100,0,100,0,4,0,0
'Access specification name,default assignment
                512K 100% Read - Sequential,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                524288,100,100,0,0,4,0,0
'Access specification name,default assignment
                512K 0% Read - Sequential,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
                524288,100,0,0,0,4,0,0
'END access specifications
Version 2006.07.27
```